Ocularis LS / Ocularis ES

RC-L / RC-E Recording Component
User Manual

0007182013040414-1457-4.0-LSES6.0

Legal Notice

This product content is intended for general information purposes only, and due care has been taken in its preparation.

Any risk arising from the use of this information rests with the recipient, and nothing herein should be construed as constituting any kind of warranty.

Disclaimer

This text is intended for general information purposes only, and due care has been taken in its preparation. Any risk arising from the use of this information rests with the recipient, and nothing herein should be construed as constituting any kind of warranty.

All names of people and organizations used in the examples in this text are fictitious. Any resemblance to any actual organization or person, living or dead, is purely coincidental and unintended.

On-Net Surveillance Systems, Inc. reserves the right to change product specifications without prior notice.

US patent # 8,390.684 B2 for Ocularis Client

Patents Applied For in the U.S. and Abroad

Table of Contents

Introductions

Product overview

This system is a fully distributed solution, designed for large multi-site and multiple server installations requiring 24/7 surveillance, with support for devices from different vendors. The solution offers centralized management of all devices, servers, and users, and empowers an extremely flexible rule system driven by schedules and events.

The Ocularis LS / ES recording component consists of the following main elements:

- The **management server** - the center of the installation
- One or more **recording servers**
- One or more **management clients,** which are unlicensed and can be downloaded and installed for free (as many times as needed).
- A **Download Manager**

Finally, your system handles an unlimited number of cameras, servers, and users—across multiple sites if required.

About updates

OnSSI regularly releases service updates for our products, offering improved functionality and support for new devices.

If you are a system administrator, OnSSI recommends that you check the <website> for updates at regular intervals in order to make sure you are using the most recent version of your system.

Management Server

What? Stores the surveillance system's camera and other configuration in a relational database, either on the management server computer itself or on a separate SQL Server on the network. Also handles user authentication, rules, etc. To enhance system performance, several management servers can be run as an OnSSI Federated Architecture (see "About OnSSI Federated Architecture" on page 212).

Where? Runs as a service, and is typically installed on a dedicated server.

What comes with the management server? When you install the management server, you get the following integrated components as well (if you select a single server management server installation (see "Install your system - Single Server option" on page 14)):

The **log server**

- **What?** Provides the necessary functionality for logging information from your system.
- **Where?** Usually installed on the same server as the management server and runs as a service.

Recording Server

What? Used for recording video and for communicating with cameras and other devices. In large installations, more than one recording server is often used on the surveillance system. Failover recording servers can be set up to take over if a recording server becomes temporarily unavailable.

Where? Recording servers as well as failover recording servers run as services, and are typically installed on separate servers rather than on the management server itself.

Management Client

What? Feature-rich administration client for configuration and day-to-day management of the system. Available in several languages.

Where? Typically installed on the surveillance system administrator's workstation or similar.

Download Manager

What? Lets surveillance system administrators manage which system-related components (e.g. particular language versions of clients) your organization's users will be able to access from a targeted web page generated by the management server. Refer to Download Manager/download web page (on page 19).

Where? Automatically installed on the management server during the installation process.

System Requirements

IMPORTANT: This system no longer supports Microsoft® Windows® XP.

For easy user/group management (see "Manage users and groups" on page 182), OnSSI recommends that you have Microsoft Active Directory® in place before you install your system. If you add the management server to the Active Directory after installing, you must re-install the management server, and replace users with new users defined in the Active Directory.

The following are *minimum* requirements for the computers used:

Computer running management server

Name	Description
CPU	Intel® Xeon®, minimum 2.0 GHz (Dual Core recommended)
RAM	Minimum 1 GB (2 GB or more recommended)
Network	Ethernet (1 Gbit recommended)
Graphics Adapter	Onboard GFX, AGP or PCI-Express, minimum 1024 x 768, 16-bit color
Hard Disk Type	E-IDE, PATA, SATA, SCSI, SAS (7200 RPM or faster)
Hard Disk Space	Minimum 50 GB free (depends on number of servers, cameras, rules, and logging settings)

Operating System	▸ Microsoft® Windows® 8 Enterprise (32 bit or 64 bit) ▸ Microsoft Windows 8 Professional (32 bit or 64 bit) ▸ Microsoft Windows 7 Ultimate (32 bit or 64 bit) ▸ Microsoft Windows 7 Enterprise (32 bit or 64 bit) ▸ Microsoft Windows 7 Professional (32 bit or 64 bit) ▸ Microsoft Windows Server 2012 (64 bit): Standard and Datacenter. ▸ Microsoft Windows Server 2008 R2 (64 bit): Standard, Web, High Performance Computing (HPC), Enterprise, and Datacenter. ▸ Microsoft Windows Server 2008 (32 or 64 bit) ▸ Microsoft Windows Server 2003 (32 or 64 bit) To run clustering/failover recording servers, you also need a Microsoft Windows Server 2003/2008 Enterprise or Data Center edition.
Software	Microsoft® .NET 3.5 SP1 and .NET 4.0 and Internet Information Services (IIS) 5.1 or newer

Computer running recording server or failover recording server

Name	Description
CPU	Dual Core Intel Xeon, minimum 2.0 GHz (Quad Core recommended)
RAM	Minimum 1 GB (2 GB or more recommended)
Network	Ethernet (1 Gbit recommended)
Graphics Adapter	Onboard GFX, AGP, or PCI-Express, minimum 1024 x 768, 16-bit color
Hard Disk Type	E-IDE, PATA, SATA, SCSI, SAS (7200 RPM or faster)
Hard Disk Space	Minimum 100 GB free (depends on number of cameras and recording settings)
Operating System	▸ Microsoft® Windows® 8 Enterprise (32 bit or 64 bit) ▸ Microsoft Windows 8 Professional (32 bit or 64 bit) ▸ Microsoft Windows 7 Ultimate (32 bit or 64 bit) ▸ Microsoft Windows 7 Enterprise (32 bit or 64 bit) ▸ Microsoft Windows 7 Professional (32 bit or 64 bit) ▸ Microsoft Windows Server 2012 (64 bit): Standard and Datacenter. ▸ Microsoft Windows Server 2008 R2 (64 bit): Standard, Web, High Performance Computing (HPC), Enterprise, and Datacenter. ▸ Microsoft Windows Server 2008 (32 or 64 bit) ▸ Microsoft Windows Vista® Business (32 or 64 bit) ▸ Microsoft Windows Vista Enterprise (32 or 64 bit) ▸ Microsoft Windows Vista Ultimate (32 or 64 bit) ▸ Microsoft Windows Server 2003 (32 or 64 bit)

Software	Microsoft® .NET 4.0 Framework.

IMPORTANT: When you format the hard disk of a recording/failover recording server device, you must change its *Allocation unit size* setting from 4 to 64 kilobytes. This is to significantly improve recording performance of the hard disk. You can read more about allocating unit sizes and find help at http://support.microsoft.com/kb/140365/en-us.

Computer running Management Client

Name	Description
CPU	Intel Core2TM™ Duo, minimum 2.0 GHz
RAM	Minimum 1 GB
Network	Ethernet (100 Mbit or higher recommended)
Graphics Adapter	AGP or PCI-Express, minimum 1024 x 768 (1280 x 1024 recommended), 16-bit color
Hard Disk Space	Minimum 100 MB free
Operating System	▸ Microsoft® Windows® 8 Professional (32 bit or 64 bit*) ▸ Microsoft Windows 8 Enterprise (32 bit or 64 bit*) ▸ Microsoft Windows 7 Professional (32 bit or 64 bit*) ▸ Microsoft Windows 7 Enterprise (32 bit or 64 bit*) ▸ Microsoft Windows 7 Ultimate (32 bit or 64 bit*) ▸ Microsoft Windows Vista® Ultimate (32 bit or 64 bit*) ▸ Microsoft Windows Vista Enterprise (32 bit or 64 bit*) ▸ Microsoft Windows Vista Business (32 bit or 64 bit*) ▸ Microsoft Windows Server 2012 (64 bit): Standard and Datacenter. ▸ Microsoft Windows Server 2008 (32 bit or 64 bit*) ▸ Microsoft Windows Server 2008 R2 (64 bit): Standard, Web, High Performance Computing (HPC), Enterprise, and Datacenter. ▸ Microsoft Windows Server 2003 (32 bit or 64 bit*) * Running as a 32 bit service/application
Software	Microsoft® .NET 4.0 Framework, DirectX 9.0 or newer, and Windows Help (WinHlp32.exe) which you can download from http://www.microsoft.com/downloads/.

Computer running log server

Name	Description
CPU	Intel® Xeon® , minimum 2.0 GHz (Dual Core recommended)
RAM	Minimum 1 GB (2 GB or more recommended)

Network	Ethernet (1 Gbit recommended)
Graphics Adapter	Onboard GFX, AGP or PCI-Express, minimum 1024 x 768, 16-bit color
Hard Disk Type	E-IDE, PATA, SATA, SCSI, SAS (7200 RPM or faster)
Hard Disk Space	Minimum 10 GB free (depends on number of servers, cameras, rules, and logging settings)
Operating System	▸ Microsoft® Windows® 8 Professional (32 bit or 64 bit) ▸ Microsoft Windows 8 Enterprise (32 bit or 64 bit) ▸ Microsoft Windows 7 Ultimate (32 bit or 64 bit) ▸ Microsoft Windows 7 Enterprise (32 bit or 64 bit) ▸ Microsoft Windows 7 Professional (32 bit or 64 bit) ▸ Microsoft Windows Server 2012 (64 bit): Standard and Datacenter. ▸ Microsoft Windows Server 2008 R2 (64 bit): Standard, Web, High Performance Computing (HPC), Enterprise, and Datacenter. ▸ Microsoft Windows Server 2008 (32 or 64 bit) ▸ Microsoft Windows Server 2003 (32 or 64 bit)
Software	Microsoft® .NET 4.0 and Internet Information Services (IIS) 5.1 or newer.

Active Directory

You normally add users from Active Directory, although you can also add users without Active Directory. Active Directory is a distributed directory service included with several Windows Server operating systems. It identifies resources on a network in order for users or applications to access them.

If you wish to add users through the Active Directory service, a server with Active Directory installed, and acting as domain controller, must be available on your network.

Installation and Removal

Installation overview

Note that the **Axis One-click Connection Component** is not installed by the common installer. This must be installed from the management server's download website (see "Download Manager/download web page" on page 19) (controlled by the Download Manager) once the management server is installed. The same applies to failover recording server(s) (see "Install failover recording server (recording server)" on page 16).

In general, your installation (or upgrade scenario) is handled by one common installer. Depending on your selections, this installer installs all or some of the following components:

- **Management Server,** the center of your system installation. Typically installed on a dedicated server.
- **Recording Server**, used for recording video feeds, communicating with cameras (via video device drivers (see "Manage video device drivers" on page 233)) and other devices. Typically installed on one or more separate computers, rather than on the machine where the management server is installed. The needed video device drivers are automatically installed along with the recording server.

  Tip: Video device drivers are small programs used for controlling/communicating with the cameras connected to a recording server. As mentioned, you get the drivers automatically during installation. However, new versions of the drivers are released from time to time and must be downloaded from the management server's download web page and installed manually.
- **Management Client**, used for configuration and day-to-day management of the system. Typically installed on the system administrator's workstation or similar.
- **Log server,** provides the necessary functionality for logging information from your system. By default installed on the management server but, if you need to increase performance, it can be installed on another server.

  When installing the log server, the URL address of the management server is expressed like this: http//123.123.123.123*. If installing the log server on the management server itself, this should be specified as* localhost*. The address can also include a port, like this:* http://123.123.123.123:2356 *or* http://Localhost:2356.

The common installer saves all components on the management server's download web page no matter whether you install them or not. Once you have run the installer, single components can be (re-)installed from the management server download web page (see "Download Manager/download web page" on page 19). Refer to Download Manager's standard installers (users) (see "Download Manager's standard installers (user)" on page 20) to see what component are available for separate download.

Since most single component installer elements are identical to the common installer elements, single component installers are not described in detail. Only exception is the **failover recording server** installer (see "Install failover recording server (recording server)" on page 16).

As well as installing on physical servers, your system installation can also take place on **virtualized servers** (see "**Install your system on virtual servers**" on page 18).

Install your system - preconditions

If you are upgrading from a previous version, refer to Upgrade from previous version (on page 25).

If you plan to run **OnSSI Federated Architecture**, refer to About OnSSI Federated Architecture (on page 212).

If you run **workgroups**, make sure to ignore the normal installation guidelines and use the alternative method for installing for workgroups indicated in the following.

- **Microsoft® Windows® Installer 4.5 - only on Windows Server 2003:**

Before installing your system, it is important to install Windows Installer 4.5**.**

- **SQL Server:**

  The management server requires access to a relational database. Later in this installation process you must choose between using an existing SQL Server on the network (**Administrator rights** on the SQL Server are required) or setting up a SQL Server Express Edition (a lighter version of a full SQL server) on the management server itself.

  If you select an SQL Server Express Edition, you might need to have Microsoft® .NET Framework 3.5 Service Pack 1 installed on the server running the SQL Server (even though Microsoft .NET Framework 4.0 is already installed). Refer to System requirements (on page 9).

- **2 x Windows Server 2003 Fix:**

  If you use Windows Server 2003, OnSSI recommends that you install two supported fixes before starting: Fix 1 and Fix 2. Otherwise, the installation of your management server might fail due to Microsoft Windows Installer process having insufficient contiguous virtual memory to verify that the *.msi* package or the *.msp* package is correctly signed.

- **Alternative installation for workgroups:**

  If you do not use a domain setup but a workgroup setup, do the following when installing:

  1. Log in to Windows using a common administrator account.
  2. Depending on your needs, start the management or recording server installation and click ***Custom***.
  3. Depending on what you selected in step 2, select to install the Management or Recording Server service using a common administrator account.
  4. Finish the installation.
  5. Repeat steps 1-4 to install any other systems you want to connect. They must all be installed using a common administrator account.

  This approach however, cannot be used when **upgrading** workgroup installations, refer to Alternative upgrade for workgroup (on page 26).

Install your system - Single Server option

In an upgrade scenario (see "Upgrade from previous version" on page 25), you might **not** want to remove the management server database as it contains your system configuration.

1. If you are installing a version downloaded from the Internet, run the `RC-L_RC-E.exe` file from the Ocularis Component Installation page. Alternatively, insert the software DVD. If the dialog does not open automatically, run the `RC-L_RC-E.exe` file from the DVD.
2. The installation files unpack. Depending on your security settings, one or more Windows security warnings may appear. Accept these and the unpacking continues. When done, the ***OnSSI*** installation dialog appears. In the coming steps, do the following:

a) Select the ***Language*** to use during the installation (this is **not** the language your system will use once installed, this is selected later). Click ***Continue***.

b) In ***Type the location of the license file***, enter the file path and name of your license file from your Ocularis provider. Preferably, use the browse function to locate it. This file has a .lic filename extension. The system verifies your license file before you can continue. Click ***Continue***.

c) Read the *OnSSI End-user License Agreement*. Select the ***I accept the terms in the license agreement*** check box. Optionally, select the ***Sign me up for the Customer Experience Improvement Program*** check box. Follow the on-screen *Read more* link for further information on this.

d) Consider the following installation methods:

- ***Single Server***, installs all management server components, recording server, and clients on the current computer. You only need to make a minimum of selections and all components are selected in the component list, which cannot be edited.
- ***Distributed***, installs all management server components and clients on the current computer. However, you must install the recording server on a separate machine. This means that the recording server is cleared in the component list which you cannot edit.
- ***Custom***, lets you select freely among management server components to install on the current computer. The only exception is the management server. By default, recording server is cleared in the component list, but you can edit this.

3. Select ***Single Server***. A list of components to install appears (you cannot edit this list). Click ***Continue***.
4. Select ***Files location*** for the program file. In ***Product language***, select the language in which your product should be installed. Click ***Install***.
5. The software now installs. When done, you see a list of successfully installed components. Click ***Close***.

   If you do not have Microsoft® IIS installed, this is automatically installed during the process. Afterwards, you may be prompted to restart your computer. Do so and after restart, depending on your security settings, one or more Windows security warnings may appear. Accept these and the installation continues. When done, your installation completes and you can get started with (see "Get started" on page 35) your surveillance system.

Install your system - Distributed option

1. Refer to Install your system - Single Server option (on page 14), steps 1-2.
2. Select ***Distributed.*** A non-editable list of components to be installed appears. Click ***Continue***.
3. Choose the type of SQL server database you want (see "Select SQL type" on page 17). Also specify the name of the SQL server. Click ***Continue***.
4. Select either ***Create new database*** or ***Use existing database*** and name the database (see "Select SQL type" on page 17). If you choose the latter, select to ***Keep*** or ***Overwrite*** existing data. Click ***Continue***.
5. Refer to Install your system - Single Server option (on page 14), step 4-5.

Install your system - Custom option

Note that with this option you can select or clear all of the components to install, except the management server. The management server is by default selected in the component list and will always be installed. If one is already installed, it will be updated.

1. Refer to Install your system - Single Server option (on page 14), steps 1-2.
2. Select ***Custom***. A list of components to be installed appears. Apart from the management server, all elements in the list are optional. The recording server is by default deselected, but you can change this if needed. Click ***Continue***.
3. Choose the type of SQL server database you want (see "Select SQL type" on page 17). If relevant, also specify the name of the SQL server. Click ***Continue***.
4. Select either ***Create new database*** or ***Use existing database*** and name the database (see "Select SQL type" on page 17). If you choose the latter, select to ***Keep*** or ***Overwrite*** existing data. Click ***Continue***.
5. Select either *This predefined account* or *This account* to select the service account (see "Select service account" on page 17). If needed, enter a password and confirm this. If you are installing a recording server and a recording server is also already installed on the same machine, this dialog is shown twice. Click ***Continue***.
6. Specify recording server properties (see "Recording/failover recording server install properties" on page 16). Click ***Continue***.

7. If you have more than one available IIS website, you can select any of these. However, if any of your websites have HTTPS binding, select one of these. Click ***Continue***.
8. Refer to Install your system - Single Server option (on page 14), step 4-5.

Install failover recording server (recording server)

IMPORTANT: During the installation process, you are asked to specify a user account under which the *Failover Server service* should run. This user account must have administrator rights in the system. Note also that if you run workgroups, you should ignore the normal installation guidelines for installing recording servers and use the alternative installation method for workgroups (see "Install your system - preconditions" on page 13).

Once you have installed the management server using the common installer, you can download the separate recording server installer from the management server's web page (see "Download Manager/download web page" on page 19) (controlled by the Download Manager). As part of this installer, you can specify whether the installer should result in a standard recording server or a failover recording server.

1. Go to the Management Server's download web page and select the Recording Server installer suitable for your needs. Save the installer somewhere appropriate and run it from here or run it directly from the web page.
2. Select the ***Language*** you want to use during the installation (this does not affect the language of your system, choose this later in the process). Click ***Continue***.
3. From a selection list of:
   - ***Typical***, which installs a standard recording server with default settings
   - ***Failover***, which installs a recording server as a failover recording server
   - ***Custom***, which installs a standard recording server and offers configuration options, for example, letting you install more than one recording server instance (see "Install multiple recording server instances" on page 25) on the current machine.

   Select ***Failover.***
4. Specify failover recording server properties (see "Recording/failover recording server install properties" on page 16). Click ***Continue***.
5. When installing a failover recording server it is mandatory to use a particular user account (*This account*) (see "Select service account" on page 17). If needed, enter a password and confirm this. Click ***Continue***.
6. Refer to Install your system - Single Server option (on page 14), step 4-5.

When the failover recording server is installed, you can check its state (see "Management Server service and Recording Server service" on page 244) from the ***Failover Server service*** icon and start using it.

More about installing

Recording/failover recording server install properties

Fill out the following properties when you install a standard recording server (see "Install your system - Custom option" on page 15) or a failover recording server (see "Install failover recording server (recording server)" on page 16):

Name	Description
Recording server name:	A name for the server in question. If required, you can later change the name.
Management server address:	The IP address (example: *123.123.123.123*) or host name (example: *ourserver*) of the management server to which the server should be connected. If required, you can later change the management server IP address/host name as part of the basic administration on the Recording server service/Failover Server service.

Media database:	The path to the media database. The media database is the recording server/failover recording server's default storage area that is the default location in which recordings from connected cameras are stored in individual camera databases. If required, you can later change the path, and/or add paths to more storage area locations.

Select SQL type

In the installer dialogs (see "Install your system - Custom option" on page 15), you must decide what to do regarding SQL database (see "Install your system - Distributed option" on page 15). The options are ***Install SQL Server 2008 Express on this computer*** / ***Use the installed Microsoft SQL Server Express database on this computer*** or ***Use an existing SQL Server on the network***. As indicated, the wording used for selecting SQL server type varies depending on whether you already have installed an SQL database on the current machine:

- First option when you have **no** SQL database installed: ***Install SQL Server 2008 Express on this computer***
  First option when you have **an** SQL database installed: ***Use the installed Microsoft SQL Server Express database on this computer***
- Second option: ***Use an existing SQL Server on the network*** is the second option.

However, it can be difficult to determine which SQL server type is right for your organization. The Microsoft SQL Server Express Edition is a "lightweight" version of a full SQL server. It is easy to install and prepare for use, and often suffices for systems with less than 300 cameras. However, if you plan to perform frequent/regular backups of your database, OnSSI recommends using an existing SQL server on the network (you must have administrator rights on the SQL server). For large installations (300 cameras or more), OnSSI recommends using a full-scale existing SQL server on a dedicated machine on the network.

IMPORTANT: OnSSI recommends that you install the database on a dedicated hard disk drive that is not used for anything else but the database. Installing the database on its own drive prevents low disk performance.

IMPORTANT: If relevant, during the database preparation process, you are asked whether you want to create a new database, use an existing database, or overwrite an existing database. For a new installation, you would typically select the default option *Create new database.* However, if you are installing the database as part of upgrading to a newer version of the system, and you want to use your existing database, make sure you select *Use existing database*.

Select service account

In the installer dialogs (see "Install your system - Custom option" on page 15), you are asked to select a service account under which the Management Server service (see "Management Server service and Recording Server service" on page 244) runs:

- With a predefined network service account (**This predefined user account**),the service always runs when the server (computer) are running - no matter which account is used. The account matters for access to various resources.
- With a particular user account (***This account***), the service uses the specified user account to run the service under the account as management server. If the server acting as management server is a member of a domain, you should either select the suggested ***Network Service*** or specify a user account for the domain in question.

> *Note that if the server in question is a failover recording server, it is **not** possible to select **This predefined account,** and when selecting* **This account***, it is **only** possible to select to specify a user account for the domain in question.*

When should I choose a particular user account instead of a predefined? If you use network drives, you should always specify a particular user account (with access to the network drives in question). Otherwise, the relevant service cannot access the required network drives.

Choose between a predefined network service account and a particular user account:

1. Select *This predefined account*.

a) Select *Network Service*.

b) Click *OK*.

- or -

1. Select *This account*.

a) Click *Browse...*. This opens the *Select User* window.

b) Verify that the relevant domain/workgroup is specified in the *From this location* field. If not, click *Locations...* to browse for the required domain/workgroup.

c) In the *Enter the object names to select* box, type the required user name. Click *OK*.

Tip: Typing part of a name is often enough. Use the *Check Names* feature to verify that the name you have entered is recognized.

d) In the *Password* field, specify the password for the user account and in the *Confirm password* field, confirm the password. The password fields cannot be empty. The password for the account must contain one or more characters and/or digits. Click *OK*.

About installer commands

As an administrator, you have a set of installer command you can use when you work with RC-L or RC-E installers.

1. On the machine where you want to enter an installer command, go to Window's ***Start***, and open a Command Prompt window.
2. In the *Command Prompt*, execute the required installer command - possible with a prefix. Note that there is a [space] before -- in all installer command lines.

Example: *RecordingServer_setup_x64.exe --ss_registration*

Tip: To get an overview of installer commands, in the *Command Prompt,* type [space]*--help* and the following window appears:

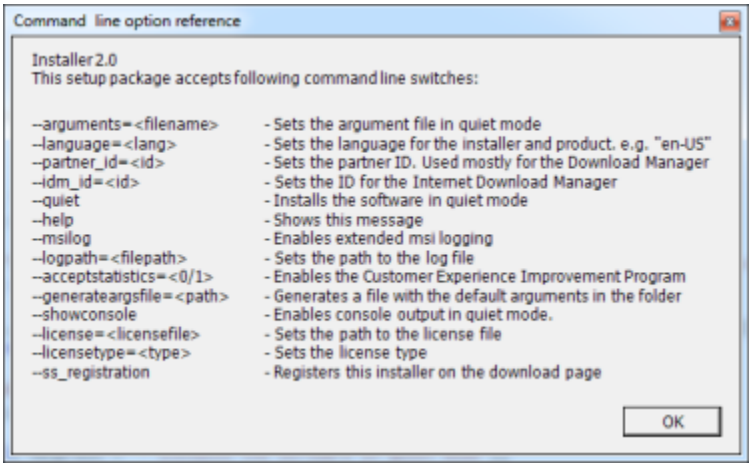

Install your system on virtual servers

You can run all system components on virtualized (see "Installation overview" on page 13) Windows® servers, such as - for example - VMware® and Microsoft® Hyper-V®. Contact your IT department for more information.

Tip: Virtualization is often preferred to better utilize hardware resources. Normally, virtual servers running on the hardware host server do not load the virtual server to a great extent, and often not at the same time. However, recording servers record all cameras and video streams. This puts high load on CPU, memory, network, and storage

system. So, when run on a virtual server, the normal gain of virtualization disappears to a large extent, since - in many cases - it will use all available resources.

If run in a virtual environment, it is important that the hardware host has the same amount of physical memory as allocated for the virtual servers and that the virtual server running the recording server is allocated enough CPU and memory - which it is not by default. Typically, the recording server needs 2-4 GB depending on configuration. Another bottleneck is network adapter allocation and hard disk performance. Consider allocating a physical network adapter on the host server of the virtual server running the recording server. This makes it easier to ensure that the network adapter is not overloaded with traffic to other virtual servers. If the network adapter is used for several virtual servers, the network traffic might result in the recording server not retrieving and recording the configured amount of images.

Download Manager/download web page

The management server has a built-in web page. This web page enables administrators and end users to download and install required RC-L / RC-E system components from any location, locally or remotely.

The web page is capable of displaying two sets of content, both by default in a language version matching the language of the system installation:

- One is targeted at **administrators**, enabling them to download and install key system components. Most often the web page is automatically loaded at the end of the management server installation and the default content is displayed. Otherwise the web page can be accessed by entering the URL:

  http://[management server address]:[port]/installation/admin/

  where [management server address] is the IP address or host name of the management server, and [port] is the port number which IIS has been set up to use on the management server. If not accessing the web page on the management server itself, log in with an account which has administrator rights on the management server.

- One targeted at end **users**, providing them access to client applications with default configuration. The content is displayed when the web page is accessed by entering the URL:

  http://[management server address]:[port]/installation/

  where [management server address] is the IP address or host name of the management server, and [port] is the port number which IIS has been set up to use on the management server.

The two web page(s) automatically has some default content so they can be used right away after the installation process. As administrator however, by using the Download Manager, you can customize what should be displayed on the web page(s). You are also able to move components between the two versions of the web page, i.e. between the one targeted at system administrators and the one targeted at end users. To move a component, right-click it, and select the web page version you want to move the component to.

Even though the Download Manager lets you control which components users can download and install, you cannot use it as a users' rights management tool. Such rights are determined by roles (see "About roles" on page 182) defined in the Management Client.

You access the Download Manager on the server running the management server software. From Windows' *Start* menu, select *All Programs*, ***OnSSI***, ***OnSSI*** *Download Manager*.

Download Manager's default configuration

As indicated, the Download Manager has a default configuration. This ensures that your organization's **users** can access standard components right from the start.

The default configuration provides **administrators a default setup** with access to downloading extra or optional components. Even though the web page, in most cases, opens automatically on the management server computer, you will often want to install key components on other servers than the management server itself. This is no problem since the web page can easily be accessed from other computers. The Download Manager's configuration is represented in a tree structure.

The first level of the tree structure simply indicates what product you are working with.

The second level refers to the two targeted versions of the web page. *Default* refers to the web page version viewed by end users. *Admin* refers to the web page version viewed by surveillance system administrators.

The third level refers to the languages in which the web page is available.

The fourth level refers to the components which are - or can be made - available to users.

The fifth level refers to particular versions of each component, which are - or can be made - available to users.

The sixth level refers to the language versions of the components which are - or can be made - available to users.

The fact that only standard components are initially available - and only in the same language version as the system itself - helps reduce installation time and save space on the server. There is no need to have a component or language version available on the server if nobody uses it.

You can, however, make more components and/or languages available (see "Add/publish Download Manager installer components" on page 20) as required. Likewise, you can hide or remove unwanted components and/or languages (see "Hide/remove Download Manager installer components" on page 21).

Download Manager's standard installers (user)

By default, the following components are available for separate installation from the management server's download web page targeted at users (controlled by the Download Manager):

- Recording servers (including failover recording servers; failover recording servers are initially downloaded and installed as recording servers, during the installation process you specify that you want a failover recording server)
- Management Client
- Log server, used for providing the necessary functionality for logging system information (see "Manage logs" on page 197)
- Axis One-click Connection Component (see "Remote connect services" on page 50) - **only available here**

For installation of **device packs**, refer to Device pack not available on Download Manager/download web page (see "Device pack installer - must be downloaded" on page 21).

Add/publish Download Manager installer components

Making non-standard components and new versions available on the management server's download page involves two procedures. First you **add new and/or non-standard components to the Download Manager**. Then you use it to **fine-tune which components should be available** in the various language versions of the web page.

If the Download Manager is open, close it before installing new components.

Adding new/non-standard files to the Download Manager:

1. On the machine where you downloaded the component(s), go to Window's ***Start***, enter a *Command Prompt*
2. In the *Command Prompt*, execute the name of the file (*.exe*) with:[space]*--ss_registration*

   Example: *RecordingServer_setup_x64.exe --ss_registration*

   The file is now added to the Download Manager, but **not** installed on the current machine.

Tip: To get an overview of installer commands, in the *Command Prompt,* type [space]*--help* and the following window appears:

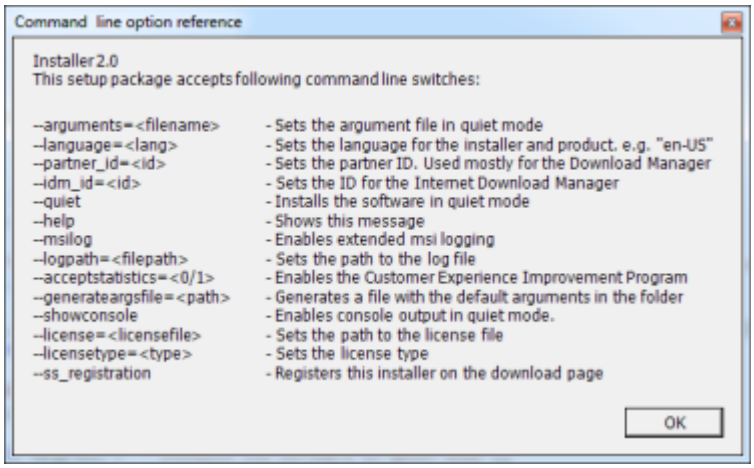

When you have installed new components they are by default selected in the Download Manager and are immediately available to users via the web page. You can always show or hide features on the web page by selecting or clearing check boxes in the Download Manager's tree structure.

Tip: You can change the sequence in which components are displayed on the web page. In the Download Manager's tree structure, drag component items and drop them at the required position.

Hide/remove Download Manager installer components

You have three options:

- **Hide components** from the web page by clearing check boxes in the Download Manager's tree structure. The components are still installed on the management server, and by selecting check boxes in the Download Manager's tree structure you can quickly make the components available again.
- **Remove the installation of components** on the management server. The components will disappear from the Download Manager, but installation files for the components are kept at *C:\Program Files (x86)\OnSSI\OnSSI Download Manager*, so you can re-install them later if required.

  1. In the Download Manager, click *Remove features...*.
  2. In the *Remove Features* window, select the feature(s) you want to remove.

     Click *OK* and *Yes*.
- **Remove installation files (see "Remove system components" on page 29) for non-required features** from the management server. This can help save disk space on the server if you know that your organization is not going to use certain features.

Device pack installer - must be downloaded

The device pack (containing device drivers (on page 233)) included in your original installation is not included on the download web page. So if you need to reinstall the device pack/make the device pack installer available, you must first add/publish the latest device pack installer to the Download Manager, by doing the following:

1. Get the newest device pack from www.onssi.com.
2. Add/publish it (see "Add/publish Download Manager installer components" on page 20) to the Download Manager by calling it with the --ss_registration (see "About installer commands" on page 18) .

Tip: If you do not have a network connection, you can reinstall the entire recording server from the Download Manager. The install files for the recording server is placed locally on your machine and in this way you automatically also get a reinstall of the device pack.

Download Manager and virus scanning

If you are using virus scanning (see "Virus scanning information" on page 249) software on the management server, it is likely that the virus scanning will use a considerable amount of system resources on scanning data from the Download Manager. If allowed in your organization, disable virus scanning on the management server.

Port numbers of special interest

Your system uses particular ports when communicating with other computers, cameras, and so on.

What is a port? A port is a logical endpoint for data traffic. Networks use different ports for different types of data traffic. Therefore, it is sometimes necessary to specify which port to use for particular data communication. Most ports are used automatically based on the types of data included in the communication. On TCP/IP networks, port numbers range from 0 to 65536, but only ports 0 to 1024 are reserved for particular purposes. For example, port 80 is used for HTTP traffic which is used when you view web pages.

In your RC-L / RC-E system, you must therefore make sure that certain ports are open for data traffic on your network.

Ports used by the system

If nothing else is stated, ports are both inbound and outbound.

- **Port 20 and 21:** Used by **recording servers** to listen for File Transfer Protocol (FTP) information; some devices use FTP for sending event messages. FTP is a standard for exchanging files across networks.
- **Port 25:** Used by **recording servers** to listen for Simple Mail Transfer Protocol (SMTP) information. Also, some devices use SMTP (e-mail) for sending event messages and /or for sending images to the surveillance system server via e-mail. SMTP is a standard for sending e-mail messages between servers.
- **Port 80:** While not directly used by the system, but by **management servers**, port 80 is typically used by the Internet Information Services (IIS) Default Web Site for running the Management Server service.
- **Port 443:** Used by the basic user authentication process where the **management server** must keep this port open at all times.
- **Port 554:** Used by **recording servers** for RTSP traffic which is used for controlling streaming from cameras.
- **Port 1024 and above** (outbound only (except ports listed in the following)): Used by **recording servers** for HTTP traffic between cameras and servers.
- **Port 5210:** Used for communication between **recording servers** and **failover recording servers** when databases are merged after a failover recording server has been running.
- **Port 5432:** Used by **recording servers** to listen for Transmission Control Protocol (TCP) information; some devices use TCP for sending event messages.
- **Port 7563:** Used by **recording servers** and **Ocularis Clients.** The main entry to the recording server where the Image Server interface is implemented. Also used for handling PTZ camera control commands and for retrieving image stream from clients etc.
- **Port 7609:** Used by the **Data Collector Server service** and must always be keep open on the machine running the **Data Collector**.
- **Port 8080**: Used for communication between internal processes on the **management server** only.
- **Port 8844:** Used for User Datagram Protocol (UDP) communication between **failover recording servers**.
- **Port 9993:** Used for communication between **recording servers** and **management servers**.
- **Port 11000:** Used by **failover recording servers** for polling (i.e. regularly checking) the state of **recording servers**.
- **Port 12345:** Used by **management servers** and **Ocularis Client** for communication.

- **Port 65101:** Used between processes on the same machine only – i.e. Inter Process Communication (IPC) on a single machine only.

Multiple management servers (cluster)

The management server can be installed on multiple servers within a cluster of servers. This ensures that the system has very little down-time. If a server in the cluster fails, another server in the cluster will automatically take over the failed server's job running the management server. The automatic process of switching over the server service to run on another server in the cluster only takes a very short time (up to 30 seconds).

Note that the allowed number of failovers is limited to two within a six hour period. If exceeded, Management Server services are not automatically started by the clustering service. The number of allowed failovers can be changed to better fit your needs. Refer to Microsoft®'s homepage for details on how to do this.

Is clustering the same as OnSSI Federated Architecture? No, clustering is not the same as federated architecture. Clustering is a method to obtain failover support for a management server on a site. With clustering, it is only possible to have one active management server per surveillance setup, but other management servers may be set up to take over in case of failure. On the other hand, federated architecture is a method to combine multiple independent sites into one large setup, offering flexibility and unlimited possibilities.

Prerequisites for clustering

- Two or more servers installed in a cluster:

  -Regarding clusters in Microsoft® Windows® 2003, refer to Deploying Microsoft® Exchange Server 2003 in a cluster.

  -Regarding clusters in Microsoft Windows 2008®, refer to Failover clusters.

- **Either** an external SQL database installed **outside** the server cluster **or** an **internal** SQL (clustered) service within the server cluster (creating an internal SQL service will require the use of SQL Server Standard or a greater version which is capable of working as a clustered SQL Server).
- A Microsoft® Windows® Server 2003/2008 Enterprise or Data Center edition.

Install in a cluster

Descriptions and illustrations might differ from what you see on your screen.

Installation and change of URL address:

1. Install the management server and all its subcomponents (see "Installation overview" on page 13) on the first server in the cluster.

   The management server must be installed with a specific user and not as a network service*. This requires that you use the* Custom *install option* ***(see "****Install your system - Custom option****" on page 15).*** *Furthermore, the specific user must have access to the shared network drive and preferably a non-expiry password.*

2. After you have installed the management server and the Management Client on the first server in the cluster, open the Management Client, *Tools,* select *Registered Services...*

a) In the *Add/Remove Registered Services* window, select *Log Service* in the list, click *Edit...*

b) In the *Edit Registered Service* window, change the URL address of the log service to the URL address of the cluster.

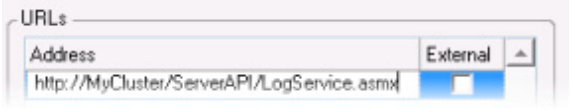

c) Repeat steps a and b for all services listed in the *Add/Remove Registered Services* window. Click *Network...*

d) In the *Network Configuration* window, change the URL address of the server to the URL address of the cluster. (This step only applies to the first server in the cluster.) Click *OK*.

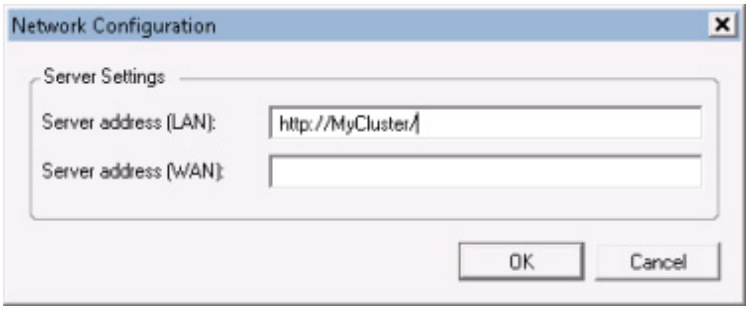

3. In the *Add/Remove Registered Services* window, click *Close*. Exit the Management Client.
4. Stop the management server service (see "Management Server service and Recording Server service" on page 244) and the IIS. Read about how to stop the IIS at Microsoft's® homepage.
5. Repeat steps 1-4 for all subsequent servers in the cluster, this time pointing to the existing SQL database. However, for the **last** server in the cluster on which you install the management server, do not stop the Management Server service.

Next, in order to take effect, the Management Server service must be configured as a generic service in the failover cluster :

1. On the last server on which you have installed the management server, go to *Start, Administrative Tools,* open Windows' *Failover Cluster Management*. In the *Failover Cluster Management* window, expand your cluster, right-click *Services and Applications*, and select *Configure a Service or Application...*

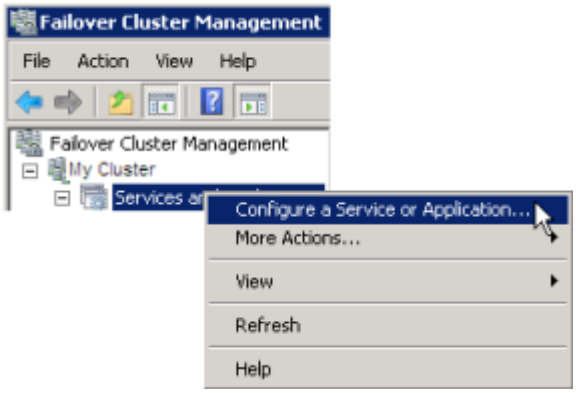

2. In the *High Availability* dialog click *Next*, select *Generic Service* and click *Next*. Do not specify anything on the third page of the dialog, click *Next*.
3. Select the *OnSSI Management Server* service, click *Next*. Specify the name (host name of the cluster) that clients use when accessing the service, click *Next***.**
4. No storage is required for the service, click *Next*. No registry settings should be replicated, click *Next***.** Verify that the cluster service is configured according to your needs, click *Next***.** The management server is now configured as a generic service in the failover cluster. Click *Finish*.

Upgrade in a cluster

Make sure to have a backup of the database in question before updating the cluster.

1. Stop the Management Server services (see "Management Server service and Recording Server service" on page 244) on all management servers in the cluster.

2. Uninstall (see "Remove system components" on page 29) the management server on all servers in the cluster.
3. Use the procedure for installing multiple management servers in a cluster as described for install in a cluster (on page 23).

> ***IMPORTANT:*** *When installing, make sure to reuse the existing SQL configuration database (which will automatically be upgraded from the old existing database version to the new one).*

Multiple recording server instances

Some information in this section may not be relevant due to differences in software versions.

It is only recommended to install multiple instances of the Recording Server service on the same server under the following conditions.

If you:

- are running Ocularis ES and are upgrading from Ocularis ES version 4.1 or older

  - and -

- are already running more 32-bit Recording Server service instances on the same server.

Since it is not possible to move devices/cameras from one recording server to another, setups running more than one 32-bit Recording Server service instances on the same server, will need to maintain this structure.

For all other setups, the newer 64-bit recording server eliminates the need for running more 32-bit instances on the same server.

Install multiple recording server instances

During the recording server installation (see "Installation overview" on page 13), select the required number of instances. A maximum of 99 recording server instances is allowed on a single server.

Using multiple recording server instances does not require additional licenses.

In the Management Client, each recording server instance will be displayed separately, allowing you to configure each instance separately.

When managing the Recording Server service (see "Management Server service and Recording Server service" on page 244) by right-clicking its icon in the notification area on the server itself, you can:

- Stop and start each instance individually
- View status messages for each instance individually, grouped on tabs.

Upgrade from previous version

This information is only relevant if you are upgrading a previous installation.

IMPORTANT: This system no longer supports Microsoft® Windows® XP (see "System Requirements" on page 9).

When upgrading, all components— **except** the management server database and if you selected so also your video device drivers—are automatically removed and replaced. The management server database is the management server's component, it contains the entire system configuration (recording server configurations, camera configurations, rules, and so on). As long as you do not remove the management server database, no reconfiguration of your system configuration is needed (although you may want to configure some of the new features in the new version).

Backward compatibility with recording servers from versions older than this current version is limited. You can still access recordings on such older recording servers, but to be able to change their configuration, they must be of the same version as this current one. Therefore, it is highly recommended to upgrade all recording servers in your system.

When you do an upgrade including your recording servers, you are asked whether you want to **update** or **keep** your video device drivers. If you choose to update, it might take a few minutes for your hardware devices to make contact with the new video device drivers after restarting your system. This is due to several internal checks being performed on the newly installed drivers.

Prerequisites

- Have your **temporary license (***.lic***) file** ready. The license file changes when your SLC changes, so you are likely to have received a new license file when you purchased the new version. When you install the management server, the wizard asks you to specify the location of your license (*.lic*) file, which the system verifies before you can continue.

  If you do not have your license file, contact your Ocularis product vendor.
- Have your **new product version** ready. If you have not purchased the software on a DVD, you can download it from www.onssi.com.
- The management server stores your system's configuration in a database. The system configuration database can be stored in two different ways:

  1. In a SQL Server Express Edition database on the management server itself

  2. In a database on an existing SQL Server on your network.

  If using 2), **Administrator rights on the SQL Server** are required whenever you want to create, move or upgrade the management server's system configuration database on the SQL Server. Once you are done creating, moving or updating, being the database owner of the management server's system configuration database on the SQL Server is sufficient.

Alternative upgrade for workgroup

If you do not use a domain setup, but a workgroup setup, you must do the following when upgrading:

1. On the recording server, create a local Windows user.
2. From the Windows ***Control Panel,*** find the ***OnSSI Data Collector service.*** Right-click it, select ***Properties***, and select the ***Log on*** tab. Set the Data Collector service to run as the local windows user you just created on the recording server.
3. On the management server, create the same local Windows user (with the same user name and password).
4. In the Management Client, add this local Windows user to the Administrator's group.

For installing with workgroups, see Install your system - preconditions (on page 13).

Installation troubleshooting

The following issues may occur during or upon installation of the management server or recording servers. For each issue, one or more solutions are available.

Issue: Recording server startup fails due to port conflict

This is an issue if the Simple Mail Transfer Protocol **(**SMTP) service is running.

It uses port 25. If port 25 is already in use, it may not be possible to start the Recording Server service. It is important that port number 25 is available for the recording server's SMTP service since many cameras are only capable of communicating via this port.

SMTP Service: Verification and solutions

To verify whether SMTP Service is installed, do the following:

1. From Windows' *Start* menu, select *Control Panel*.
2. In the *Control Panel*, double-click *Add or Remove Programs*.
3. In the left side of the *Add or Remove Programs* window, click *Add/Remove Windows Components*.
4. In the *Windows Components* wizard, select *Internet Information Services (IIS)*, and click *Details...*.
5. In the *Internet Information Services (IIS)* window, verify whether the *SMTP Service* check box is selected. If so, SMTP Service is installed.

If SMTP Service is installed, select one of the following solutions:

Solution 1: Disable SMTP Service, or set it to manual startup

This solution lets you start the recording server without having to stop the SMTP Service every time:

1. From Windows' *Start* menu, select *Control Panel*.
2. In the *Control Panel*, double-click *Administrative Tools*.
3. In the *Administrative Tools* window, double-click *Services*.
4. In the *Services* window, double-click *Simple Mail Transfer Protocol (SMTP)*.
5. In the *SMTP Properties* window, click *Stop*, then set *Startup type* to either *Manual* or *Disabled*.

   Tip: When set to *Manual*, the SMTP Service can be started manually from the *Services* window, or from a command prompt using the command *net start SMTPSVC*.
6. Click *OK*.

Solution 2: Remove SMTP service

Note that removing the SMTP Service may affect other applications using the SMTP Service.

1. From Windows' *Start* menu, select *Control Panel*.
2. In the *Control Panel*, double-click *Add or Remove Programs*.
3. In the left side of the *Add or Remove Programs* window, click *Add/Remove Windows Components*.
4. In the *Windows Components* wizard, select the *Internet Information Services (IIS)* item, and click ***Details...***.
5. In the *Internet Information Services (IIS)* window, clear the *SMTP Service* check box.
6. Click *OK, Next*, and *Finish*.

Issue: Manual installation of IIS if needed

When installing, the Internet Information Services (IIS) is under normal circumstances automatically installed. If the automatic installation fails, you must install the IIS manually:

1. If automatic IIS installation fails, you will see an error message asking you to install the IIS manually. In the error message box, click *Install IIS Manually*.
2. Select *Server Manager* from Windows' *Start* menu. In the left side of the *Server Manager* window, select *Roles*, then the *Roles Summary*.
3. Now select *Add Roles* to start a wizard.

4. In the wizard, click *Next*, select *Web Server (IIS)*, and follow the wizard's steps.
5. When you reach the wizard's ***Select Role Services*** step, you will see that some role services are selected by default. However you should select some additional role services:
   - Under *Security*, select *Basic Authentication and Windows authentication*.
   - Under *Management Tools*, select *IIS Management Console*, expand it, and select *IIS 6 Metabase Compatibility, IIS 6 WMI Compatibility, IIS 6 Scripting Tools*, and *IIS 6 Management Console***.**

     When ready, the relevant part of the *Role services* tree should look like this:

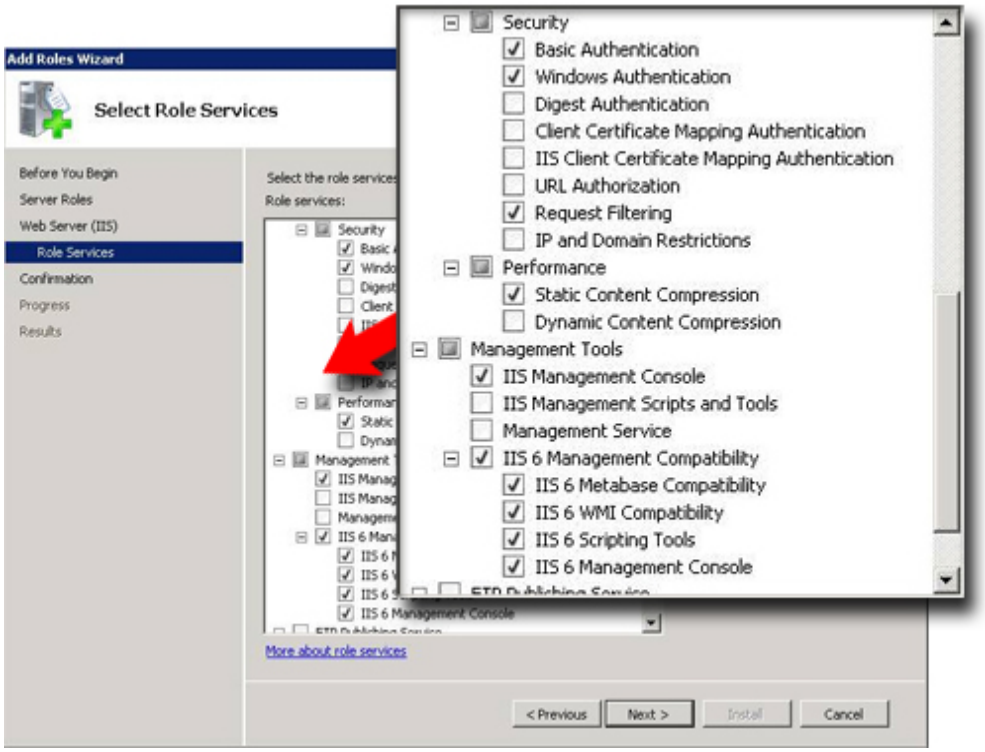

6. Complete the wizard by following the remaining steps.

Issue: Changes to SQL server location prevents database access

This is an issue if the location of the SQL Server is changed, for example by changing the host name of the computer running the SQL Server. The result of this issue will be that the access to the database is lost.

Solution: Use the update SQL address tool (see "Update SQL server address" on page 243)found at the tray icon, aka Systray.

Issue: Insufficient continuous virtual memory fails installation

The following is only relevant if you use **Windows Server 2003.**

If you try to install a large Windows Installer package or patch package in Windows Server 2003, this problem might occur if the Windows Installer process has insufficient continuous virtual memory to verify that the *.msi* package or the *.msp* package is correctly signed.

Solution: A supported fix is available for Windows Server 2003.

Issue: Multi-domain environments; one-way trusts not working

Refer to Setup with one-way trust (on page 254).

Remove system components

The following procedure describes standard system component removal in recent Windows versions; the procedure may be slightly different in older Windows versions:

1. In Windows' *Start* menu, select *Control Panel*, and then...
   - If using *Category* view, find the *Programs* category, and click *Uninstall* a program.
   - If using *Small icons* or *Large icons* view, select *Programs and Features*.
2. In the list of currently installed programs, right-click the required OnSSI program or service.
3. Select *Uninstall* if you wish to uninstall all components. Select *Change* if you only wish to uninstall some components
4. Follow the removal instructions.

Remove recording server

To remove a **recording server** installed on another machine than the management server, use the following procedure on the computer on which the recording server is installed:

1. Stop the Recording Server service by right-clicking the recording server icon in the computer's notification area (also known as the *system tray*), then select *Stop Recording Server service*.

Recording server notification area icon

2. To remove, follow the general removal procedure (see "Remove system components" on page 29).
3. Right-click the *Recording Server* in step 2 of the general removal process.

What happens to the recording server's recordings? During the removal process, you are asked whether you want to keep the recording server's recordings.

Management Client

Management Client Overview

The Management Client is the administration client used for configuration and day-to-day administration of the recording component (Management Server). The Management Client software is typically installed (see "Installation overview" on page 13) on the surveillance system administrator's workstation or similar.

Management Client's elements

Depending on the recording component, functionality described here may be limited or unavailable.

The Management Client window is divided into a number of panes. The number of panes will change depending on your task:

The following illustrations outline the Management Client window's default layout; the window layout can be customized (see "Customize the Management Client's layout" on page 39), and may therefore be different on your computer.

- When working with recording servers and devices (cameras, inputs, outputs), the Management Client window contains a menu bar and four panes (see "Panes Overview" on page 33):

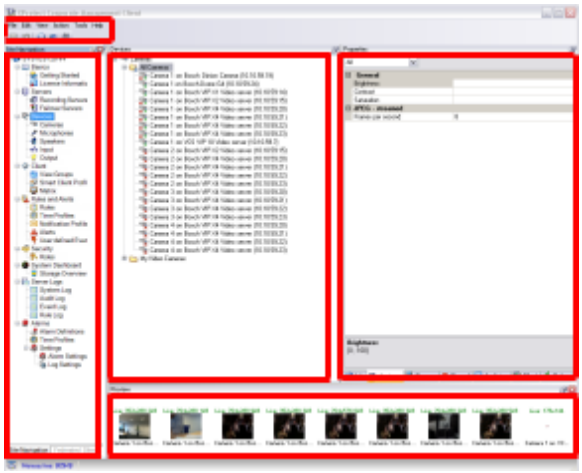

- When working with rules, time and notification profiles, users, roles, and so on, the Management Client window typically contains a menu bar and three panes:

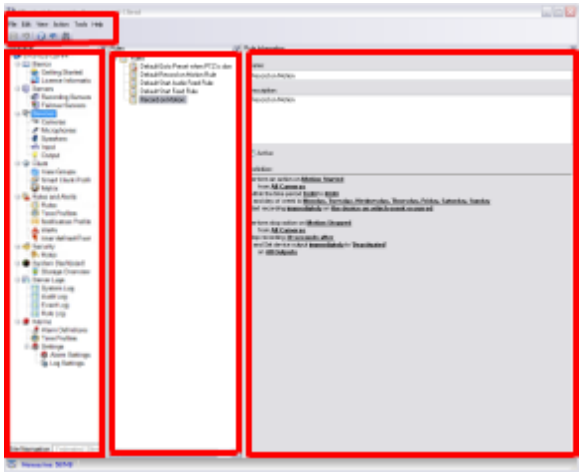

- When viewing logs, the Management Client window typically contains only the menu bar, the Site Navigation and Federated Sites Hierarchy Pane (see "Panes Overview" on page 33) and an overview area (marked in gray):

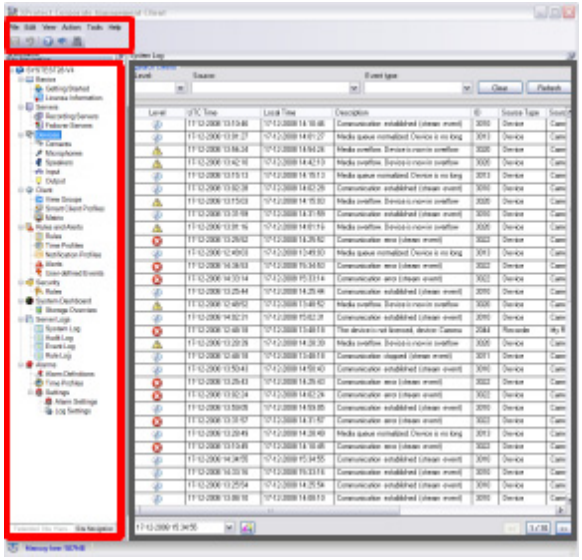

Site Navigation pane and Federated Hierarchy pane

The Management Client is the feature-rich administration client used for configuration and day-to-day administration of your system. The Management Client software is typically installed on the surveillance system administrator's workstation or similar.

Site Navigation pane: Your main navigation element in the Management Client. Name, settings and configurations of the site you are logged into are reflected (see "Manage OnSSI Federated Architecture" on page 219) here (site-name is visible at the top of the pane). The Management Client's features are grouped into categories reflecting the functionality of the software.

Tip: Right-clicking items in the Site Navigation pane gives you quick access to management features.

Federated Site Hierarchy pane: Your navigation element dedicated to displaying OnSSI Federated Architecture sites (see "About OnSSI Federated Architecture" on page 212) and their parent/child links.

The parent server you are logged in to, your home site, is always at the top, and adopting its point of view, you can view all its linked children and downwards in the parent/child hierarchy.

What if I only have one server and don't run OnSSI Federated Architecture? Your user interface looks the same, but you only see the one server in your setup.

Menu Bar

The Management Client's menu bar features the following menus (see "Management Client Menu Overview" on page 37):

File, *Edit*, *View*, *Action*, *Tools* and *Help*.

Toolbar

The Management Client's toolbar features the following options:

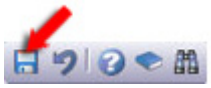

Save: Save changes to your settings.

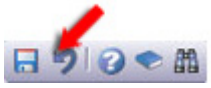

Undo: Undo your latest change.

Help...: Access a help topic relevant to your task

Contents...: Access the help system's table of contents.

Search...: Access the help system's search feature.

Memory Indicator

The memory indicator in the lower left corner of the Management Client states how much memory is available for working with the Management Client.

When you expand items in the Site Navigation pane (see "Panes Overview" on page 33), the Management Client uses memory to treat data stored in the individual items. Expanded items keep processing even when you expand other items, letting you access already-expanded items faster.

When available memory drops to 300 MB the memory indicator numbers turn red:

When the memory indicator drops to 0 MB, you cannot expand any more items and will see a warning dialog asking you to free up memory by refreshing your configuration.

Click *OK* to exit this dialog and press **F5** on your keyboard or select *Refresh* in the *Action* menu.

Panes Overview

Depending on the recording component, functionality described here may be limited or unavailable.

The Management Client contain the following panes:

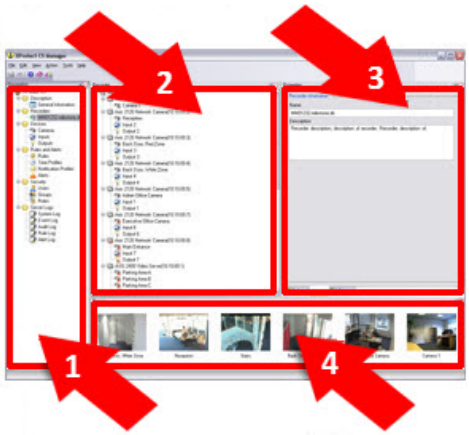

1. Site Navigation Pane and Federated Sites Hierarchy Pane
2. Overview pane
3. Properties pane
4. Preview pane

The illustration outlines the Management Client window's default layout; the window layout can be customized (see "Customize the Management Client's layout" on page 39), and may therefore be different on your computer.

Menu and tool bars: Provide quick access to often-used features.

Site Navigation pane: Your main navigation element in the Management Client. Name, settings and configurations of the site you are logged into are reflected (see "Manage OnSSI Federated Architecture" on page 219) here (site-name is visible at the top of the pane). The Management Client's features are grouped into categories reflecting the functionality of the software.

Tip: Right-clicking items in the Site Navigation pane gives you quick access to management features.

Federated Site Hierarchy pane: Your navigation element dedicated to displaying OnSSI Federated Architecture sites (see "About OnSSI Federated Architecture" on page 212) and their parent/child links.

The parent server you are logged in to, your home site, is always at the top, and adopting its point of view, you can view all its linked children and downwards in the parent/child hierarchy.

What if I only have one server and don't run OnSSI Federated Architecture? Your user interface looks the same, but you only see the one server in your setup.

Overview Pane: Provides overview of the item you have selected in the Site Navigation Pane, typically in the form of a detailed list. Selecting a particular item in the Overview pane will typically display the item's properties in the Properties pane. Right-clicking items in the Overview pane gives you access to management features.

Properties pane: Displays properties of the item selected in the Overview pane. In many cases, properties are displayed across a number of tabs:

Example of properties displayed on tabs

Preview pane: You will see the Preview pane when you deal with recording servers and devices. It displays preview images from selected cameras or states information from selected microphones, speakers, inputs and outputs. The example shows a camera preview image with information about the resolution and data rate of the camera's live stream:

By default, information shown with camera preview images will concern live streams (shown in green text). If you want recording stream information instead (shown in red text), in the Management Client's menu, select ***View*** > ***Show Recording Streams***.

Toggle the Preview pane on and off in the ***View*** menu. To resize the Preview pane, drag its borders. The larger the Preview pane, the larger preview images and state information will appear.

Performance can be affected if the Preview pane displays preview images from many cameras at a high frame rate. To control the number of preview images, and their frame rate, in the Tools menu, select ***Options*** > ***General***.

Basics

Get started

Here the tasks typically involved in setting up the system are listed.

Note that although information is presented as a checklist, a completed checklist does not in itself guarantee that the system will match the exact requirements of your organization. To make the system match the needs of your organization, it is highly recommended that you monitor and adjust the system once it is running.

For example, it is often a good idea to spend time on testing and adjusting the motion detection sensitivity settings of individual cameras under different physical conditions (day/night, windy calm weather, and so on) once the system is running. The setup of rules, which determine most of the actions performed by the system (including when to record video), is another example of configuration which to a very large extent depends on your organization's needs.

☑ Install (see "Installation overview" on page 13) the various components of your system.

☑ Log in (see "Use the Management Client to Log in to the Management Server" on page 37) to the Management Client

☐ Authorize use (see "Authorize a recording server" on page 66) of your system's recording servers.

Why must I authorize recording servers? In a surveillance system, recording servers point to management servers, not the other way round. In theory, recording servers which you do not want to include in your surveillance system could thus be configured to connect to your management servers. By authorizing recording servers before they can be used, surveillance system administrators have full control over which recording servers are able to send information to which management servers.

☐ Detect the hardware devices (see "Add hardware" on page 53) (for example, cameras and video encoders) which should be added to each recording server.

What is the Add Hardware wizard? *Add Hardware* helps you detect IP hardware devices, such as cameras and video encoders, on your network and add them to your system. The wizard offers you two ways of detecting and adding hardware devices: With *automatic hardware detection*, the system automatically scans for available hardware within one or more specified IP address ranges. With *assisted hardware detection*, you manually specify the IP address of each required device. Both options offer the possibility of automatically detecting the correct hardware drivers.

☐ Verify that each recording server's storage areas will meet your needs (see "About storage and archiving" on page 61)

What is a storage area? A storage area is a directory in which the databases containing recordings from the cameras connected to the recording server are stored— each individual camera database by default has a maximum size of 5 GB. A default storage area is automatically created for each recording server when the recording server is installed on the system. Connected cameras' databases are stored in the recording server's default storage area unless you specifically define that another storage area should be used for storing the databases of particular cameras. If required, a wizard lets you add further storage areas (on the recording server computer itself, or at another location, for example on a network drive), edit which storage area should be the default area, and so on

☐ Verify that each recording server's archiving settings will meet your needs (see "About storage and archiving" on page 61).

What is archiving? Archiving is the automatic transfer of recordings from a camera's default database to another location. This way, the amount of recordings you are able to store will not be limited by the size of the camera's default database. Archiving also makes it possible to back up your recordings on backup media of your choice. Archiving is configured on a per-recording server basis. Once you have configured the archiving settings for a recording server (where to store archives, how often to transfer recordings to the archives, and so on), you can enable archiving for individual cameras. When archiving is enabled for a camera, the contents of the camera's database will automatically be moved to an archive at regular intervals.

☐ Configure any required failover recording servers (see "About failover recording servers—regular and hot standby" on page 234). A failover recording server is a spare recording server which can take over if a standard recording server becomes unavailable.

☐ Configure each recording server's individual cameras (see "Manage cameras" on page 82).

Tip: You are able to group cameras, and configure common properties for all cameras within a group in one go.

Tip: Motion detection, a vital setting on most IP surveillance systems, is enabled by default. However, you may want to fine-tune motion detection settings, or disable motion detection for particular cameras.

☐ Enable and configure microphones (see "Manage Microphones" on page 102)— if any.

☐ Enable and configure speakers (see "Manage speakers" on page 107)— if any.

☐ Enable and configure input (see "Manage input" on page 109)— if any.

☐ Enable and configure output (see "Manage output" on page 113)— if any.

☐ Create rules (see "Manage rules" on page 165).

What is a rule? Rules are a central element in your system. The behavior of the system is to a very large extent determined by rules. Rules determine highly important settings, such as when cameras should record, when PTZ (Pan/Tilt/Zoom) cameras should patrol, when notifications should be sent, and so on

Tip: When creating rules, you may also want to use time profiles (see "Manage time profiles" on page 173) (for quickly making rules apply within or outside predefined periods of time) or notification profiles (see "Manage notification profiles" on page 176) (for quickly making rules send preconfigured e-mails— with video clips, if required— to selected recipients).

☐ Add roles (see "Manage roles" on page 184).

What is a role? Roles determine which system features users and groups are able to use. In other words, roles determine rights.

☐ Add users and/or groups of users (see "Manage users and groups" on page 182).

Tip: If you have a server with Active Directory installed, and acting as domain controller on your network, the system lets you quickly add users and/or groups from Active Directory. Only one user is necessary when using the recording component with Ocularis. A second user may be required if also using Ocularis OpenSight.

☐ Activate licenses (see "About licenses" on page 47).

Why must licenses be activated? When installing the system, you used a single temporary license. The temporary license is only valid for a certain number of days. After this initial period ends, all recording servers and cameras on your system will require activation of their individual licenses. You must therefore activate your licenses before the initial period ends, since all recording servers and cameras for which no licenses have been activated will otherwise stop sending data to the surveillance system.

Use the Management Client to Log in to the Management Server

Access to the Management Client requires certain user rights. Consult your surveillance system administrator if in doubt.

1. Click the Management Client desktop icon or—in Windows' *Start* menu—select *All Programs* > *OnSSI Management Client*. This makes the login window appear.
2. In the *Computer* field, type the name of the computer running the management server (leaving out http/https in front).
3. You have three different options when logging in: ***Windows Authentication (current user)***, ***Windows Authentication***, and ***Basic Authentication***.
4. By default, you log in with your active Windows account. This means that if you are currently logged in as, for example, *JohnSmith*, by default you log in to the management server as *JohnSmith* as well.
5. Depending on how you wish to log into the management server, in the *Authentication* field select:
   - *Windows Authentication (current user)* if you want to log in with your active Windows account (this is the default login option).
   - *Windows Authentication*, if you want to log in with a different Windows account.
   - Basic Authentication, if you want to log in with a basic user authentication.

   For ***Windows Authentication*** and ***Basic Authentication*** also fill in the *User name* and *Password* fields respectively.

   Tip: If you have logged in with a specific user type before (*Windows Authentication*, *Basic Authentication,* or both) you can select previously entered user names in the user name list.
6. Click *Connect* to open the Management Client software.

Management Client Menu Overview

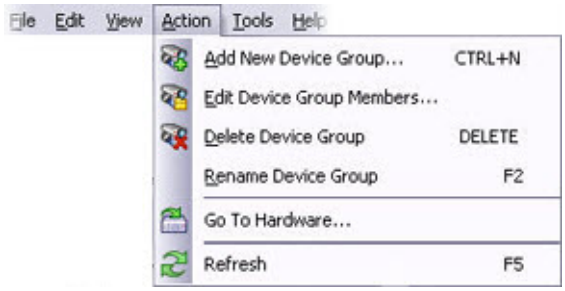

Example only; Some menus may change depending on context.

Action menu items

(Depending on context)

Name	Description
Refresh	Is always available and reloads the requested information from the management server.

Expand (or *Collapse*)	Is available when working with *Federated architecture*, *Servers*, *Devices, Client, Rules* and *Events* and *System Dashboard*.
A number of context specific items	If relevant.

Be aware of the following when working with the *Action* menu concerning OnSSI Federated Architecture (see "About OnSSI Federated Architecture" on page 212). To be able to delete a site without being connected to it (see "Manage OnSSI Federated Architecture" on page 219), **right-clicking a site does not select it, but offers a context menu**. Because of this, some context menu items may be disabled if you are not connected to the site and some are only available on the home-site, i.e. the site you are logged in to.

Edit menu items

Name	Description
Undo	Cancel your latest action.

File menu items

Name	Description
Save	Save your current configuration.
Logoff...	Log out of the Management Client, and log in with another user account if necessary.
Exit	Close down and exit the Management Client.

Help menu items

Name	Description
Help...	Access a help topic relevant to your task.
Contents...	Access the help system's table of contents.
Search...	Access the help system's search feature.
About...	Opens a dialog displaying information about the version of your Management Client.

Tools menu items

Name	Description
Registered Services...	Add registered servers.
Ocularis CS Servers...	Add Ocularis CS servers (see "Manage Ocularis CS servers" on page 202) specifically. Only relevant if you run Ocularis ES.

Effective Roles...	View all roles of a selected user or group (see "Manage users and groups" on page 182). Only relevant if you run Ocularis ES.
Options...	Opens the Options dialog (see "Options" on page 207), which lets you define and edit several global system settings. Only relevant if you run Ocularis ES.

View menu items

(Depending on context)

Name	Description
Reset Application Layout	Reset the layout (see "Customize the Management Client's layout" on page 39) of the different panes in the Management Client to their default settings.
Preview Window	Toggle the Preview pane (see "Panes Overview" on page 33) on and off when working with recording servers and devices. **Tip:** If the ***Preview*** pane displays images from many cameras at a high frame rate, it may slow down performance. To specify the number of preview images you want in your ***Preview*** pane, as well as their frame rate, select *Options* > *General* from the *Tools* menu.
Show Recording Streams	By default, the information shown with preview images in the Preview pane will concern cameras' live streams (shown in **green** text). If you want information about recording streams instead, select *Show Recording Streams*. Recording stream information will be shown in **red** text.
Federated Site Hierarchy	By default, the Federated Site Hierarchy pane is enabled, and this command lets you toggle it on and off.
Site Navigation	By default, the Site Navigation pane (see "Panes Overview" on page 33) is enabled, and this command lets you toggle it on and off.

Customize the Management Client's layout

You can rearrange panes in the Management Client, and customize its look to suit your needs. If you rearrange the panes, you can always reset the entire layout to the Management Client's default layout.

Resize panes

You can resize panes by dragging the borders of the panes:

1. Place your mouse pointer over a border.
2. When the pointer becomes a double-headed arrow, drag the border in the required direction.

The size of the content inside the panes stays the same regardless of the size of the panes, with one exception: the larger the Preview pane (see "Panes Overview" on page 33) is, the larger preview images and state information will appear.

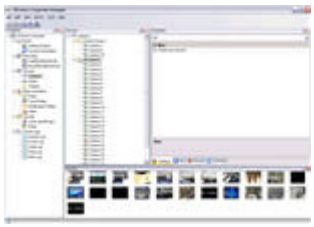

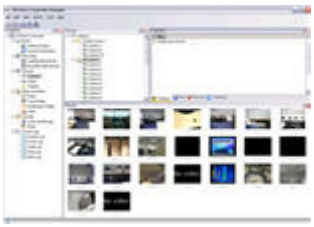

Move panes

You can move a pane to a different position either as a floating pane or to a docked position, by clicking on a pane's title bar and dragging it with the mouse.

The position and whether the pane becomes a floating pane or docked depend on where you release the mouse button.

The Management Client offers some layout elements that help you control the new position of the pane. The layout elements a

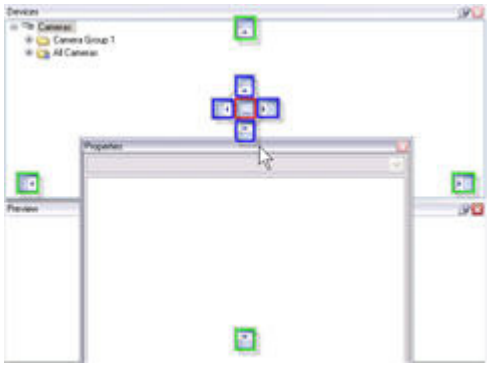

Outer lay elements illustrated with **green,** inner layouts with **blue** and center layouts with **red**

Floating panes

To move a pane to a floating pane, drag the pane to its new position *without* using one of the layout elements.

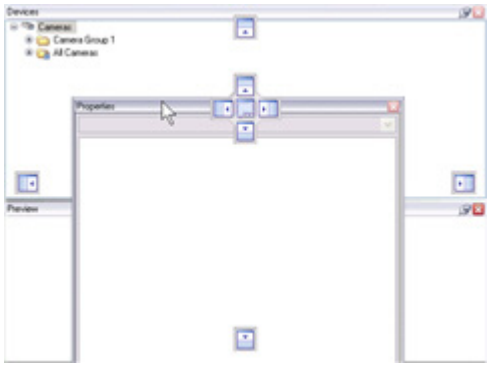

Dragging a pane to a position without using a layout element

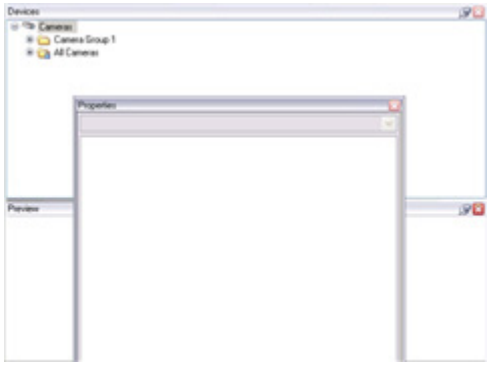

Result: A floating pane

Move a pane to a docked outer position

If you move a pane to a docked outer position, it fills the area with a horizontal or vertical split that goes from top to bottom or left to right.

1. Drag the pane to one of the outer layout elements.

   Tip: Before you release the mouse, the pane's new position is indicated by a gray area.

2. Release the mouse to dock the pane at its current position.

Dragging a pane to the right outer layout element

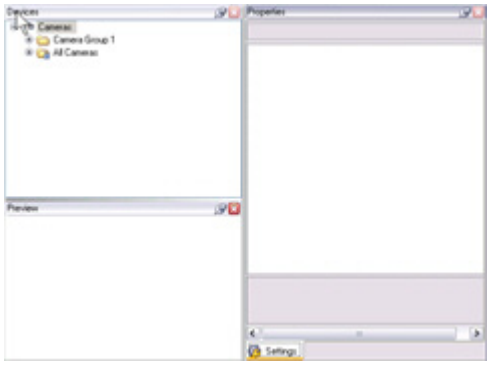

Result: The pane is docked to the right

Move a pane to a docked inner position

If you drag the pane to one of the inner layout elements, the pane will be positioned along one side of one of the other panes.

1. Drag the pane to one of the inner layout elements.

Tip: Before you release the mouse, the pane's new position is indicated by a gray area.

2. Release the mouse to dock the pane at its current position.

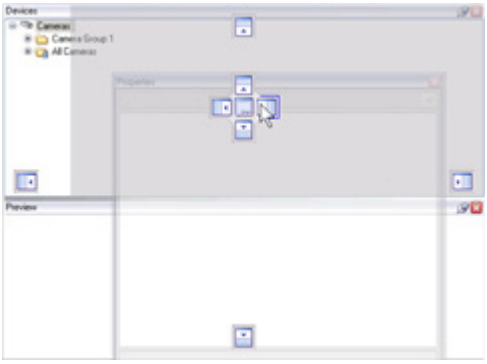

Dragging a pane to the right inner layout element of the Overview pane

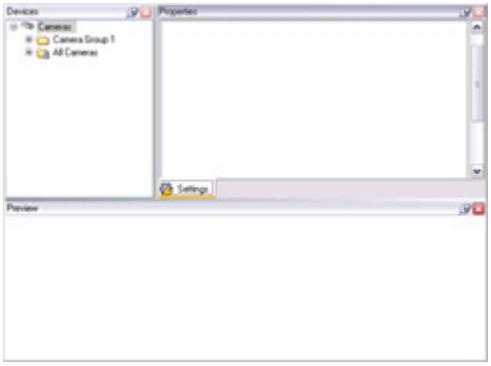

Result: The pane is docked to the right of the Overview pane

Move a pane to a shared position

You can move a pane into another pane's position so two or more panes share the same position:

1. Drag the pane to the center layout element of the pane which position you want to share.

The center layout element

Tip: Before you release the mouse, the pane's new position is indicated by a gray area.

2. Release the mouse to dock the pane at its current position.

Tip: To view the content of the panes, click the tabs on the bottom of the shared position.

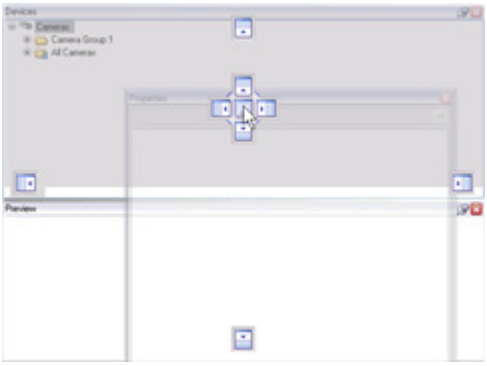

Dragging a pane to the inner center layout element of another pane

Result: The pane shares the same position as the other pane

Split shared positions

If you do not want a pane to share a position with another pane, do this:

1. Click the tab of the relevant pane and drag it to a new position.

   The pane's new position can be a docked position or a floating pane.
2. Release the mouse to place the pane at its current position.

Use auto-hide

You can auto-hide panes. An auto-hidden pane is available as a tab to the right or left of the previous position of the pane. When you place your mouse pointer over the tab, the content of the pane slides out. As soon the cursor is positioned outside the pane, it slides back.

To auto-hide a pane click the *Auto Hide* pushpin in the title bar of the pane you want to auto-hide.

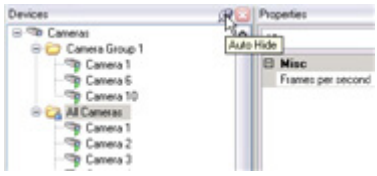

Default appearance and position of the Overview pane.

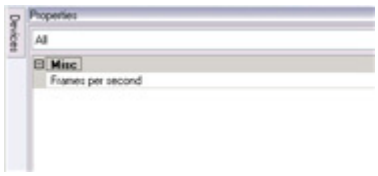

The Overview pane is hidden and available through a tab to the left.

Do the following to show and open an auto-hidden pane again:

1. Place your mouse pointer over the tab of the auto-hidden pane to show the pane.
2. Click the *Auto Hide* pushpin in the title bar of the pane to dock the pane.

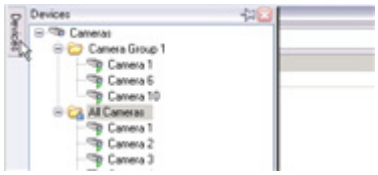

Reset to default layout

If you have moved, resized and auto-hidden panes and now want to reset the entire layout of the panes in the Management Client to their default settings, do the following:

1. From the Management Client's *View* menu, select *Reset Application Layout*.
2. Restart the application.

Toggle Preview pane on and off

You can close the Preview pane (see "Panes Overview" on page 33) when working with recorders and devices by clicking *Close* in the right side of the Preview pane's title bar.

To reopen the Preview pane select *Preview Window* from the Management Client's *View* menu.

Tip: If the Preview pane displays images from many cameras at a high frame rate, it may slow down performance. To specify the number of preview images you want in your Preview pane, as well as their frame rate, select *Options* > *General* from the *Tools* menu.

Tip: When the Preview pane is closed, it uses no resources and improves therefore the computer's performance.

Activate (Register) Licenses - Online or Offline

Once cameras have been added to the recording component using the Management Client, the camera is fully functional. You may configure its settings and start to use it. A timestamp of when the camera was installed will be recorded and the camera will work normally for a 30 day grace period. You have 30 days to register the camera licenses.

Registering camera licenses may be done online or offline. If the computer where the Management Client is installed has connectivity to the internet, use the online method. If not, use the offline method.

Keep in mind that if there are multiple recording servers in the system, cameras on each recorder must go through this process and be registered.

You can view the status of the recorder's camera licenses from the Site Navigation Pane. Select License Information from the Basics node.

Activate (Register) Licenses - Online

1. You must first register with the ***OnSSI Licensing Portal***. This process need only be done once. Open a browser and go to the following URL:

   https://rclicensing.onssi.com/CustomerRegistration

2. Enter your email address. This will become your user name for the OnSSI Licensing Portal.

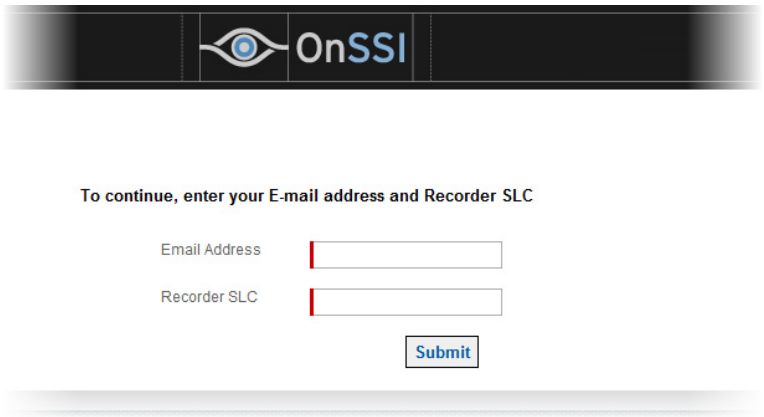

OnSSI License Portal Page

3. Enter the SLC for the recorder. This is the alphanumeric value found on the license file. For instance, the code should have 11 characters and be in the form similar to: XXX-XXXX-XXXX.
4. Click Submit.
5. You will receive an email with your OnSSI Licensing Portal user name and password.
6. In the Management Client, right-click *License Information* and select *Activate License Online*.

   An Activate Online pop-up appears.

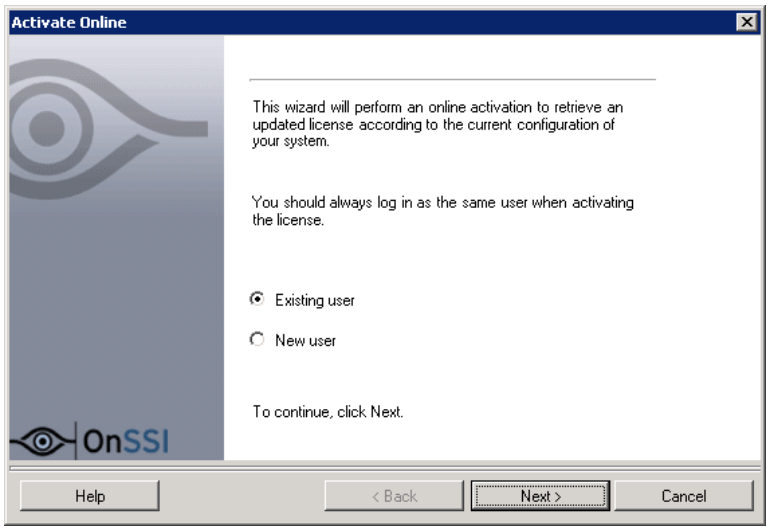

Select Existing User

7. Select the *Existing user* radio button.

   It is important to select Existing user. DO NOT select New User.

8. Click **Next**.

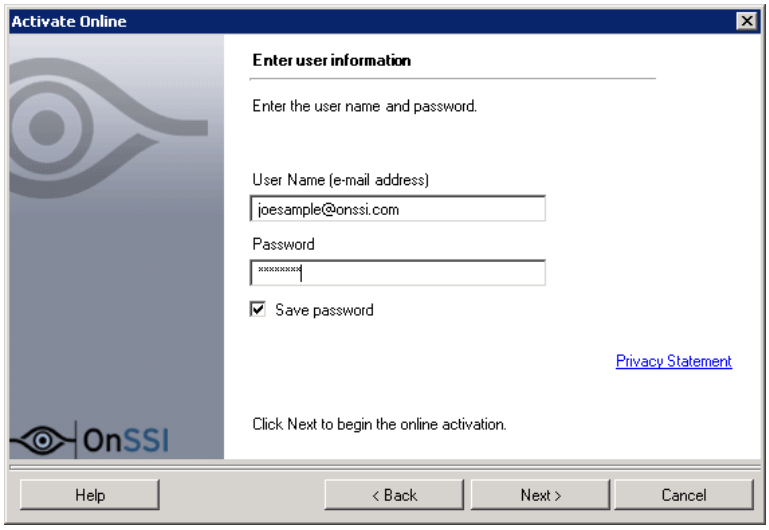

Enter OnSSI License Portal credentials

9. Enter the user name (email address) and password that was emailed to you in step 2. Click **Next**.

Once camera licenses have been validated, the camera license count will go from *Temporary* to *Activated* in the License Information screen.

If at any point in this license registration process you have questions or receive an error message, please contact OnSSI Sales.

Activate (Register) Licenses - Offline

Use this procedure when the management client computer does not have internet connectivity.

1. On the Management Client computer, open the Management Client and right-click *License Information* and select *Activate License Offline > Export License For Activation.*

   A Save Request File pop-up appears. The software will gather the necessary system data and package in a license request file with a .lrq filename extension. The default name for the file is the SLC. Use this default name and store the file in an accessible location.

2. Click Save.
3. Copy the .lrq file to portable media.
4. On any computer with internet connectivity, email the .lrq file as an attachment and send to support@onssi.com.

   OnSSI will process the file and return to you an email with a new license file (.lic) as an attachment.

5. Detach this file to portable media and bring to management client computer.
6. On the Management Client computer, open the Management Client and right-click *License Information and select Activate License Offline > Import Activated License*.
7. Locate and select the new .lic file.
8. Click Open.

Once camera licenses have been validated, the camera license count will go from *Temporary* to *Activated* in the License Information screen.

If at any point in this license registration process you have questions or receive an error message, please contact OnSSI Sales.

Activate licenses after grace day period

If the grace day period is exceeded before activation, all cameras which are not activated within the given period will become unavailable, and will not be able to send data to the surveillance system.

If you exceed the grace day period before you activate a license, the license is not lost. You can activate the license as usual. Configuration, added cameras, defined recording servers, and other settings will not be removed from the Management Client if a license is activated too late.

About licenses

When you purchase the system, you also purchase a certain number of licenses for device channels. Device channels are typically cameras but could also be dedicated input/output boxes.

At first, when you have installed the various system components, configured the system, and added recording servers and cameras through the Management Client, the surveillance system runs on temporary licenses which need to be activated before a certain period ends. This is the grace day period.

When the new surveillance system is working, we recommend that you activate your licenses (see "Activate (Register) Licenses - Online or Offline" on page 44) before you make the final adjustments. The reason is that you must activate your licenses before the grace day period expires, since all recording servers and cameras for which no licenses have been activated will not be able to send data to the surveillance system if the grace day period is expired.

Devices which require a license

You need licenses for the number of device channels you want to run on the system. Device channels are typically cameras but could also be dedicated input/out boxes. One device channel license enables you to run one camera or one dedicated input/output box. You can use and define an unlimited number of recording servers, microphones, speakers, inputs and outputs.

You can always get more licenses (see "Get additional licenses" on page 48) as your surveillance system grows.

License information

To get an overview of licenses in your system, go to the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Basics*, and select *License Information*. This brings up the *License Information* page displaying:

- your product type
- your software license code
- the number of available device channels (typically cameras, but it could also be dedicated input/output boxes) and cameras.
- if you run OnSSI Interconnect (see "About OnSSI Interconnect" on page 57), the total number of available OnSSI Interconnect device channels and OnSSI Interconnect cameras.
- how many licenses you have used, both the number of activated licenses and the number of temporary (not activated) licenses
- whether you need to get additional licenses in order to have enough licenses for all of your cameras, the number of additional license you need, is found by adding your missing licenses with your expired licenses
- any other installed products used with the system, and—if applicable—their Software License Code.

Note that the *License Information* page does **not** list microphones, speakers and in- and output boxes since these are unlimited.

In the Site Navigation pane you can activate licenses (see "Activate (Register) Licenses - Online or Offline" on page 44) by expanding *Basics* and right-clicking *License Information*.

The cameras for which you do not have a license will not send data to the surveillance system. Cameras added after all available licenses are used are unavailable. Cameras without licenses will be identified by an exclamation mark symbol when listed in the Management Client's Overview pane (see "Panes Overview" on page 33).

Tip: In the short period until you have obtained additional licenses, you can disable some less important cameras to allow some of the new cameras to run instead.

Refer to Manage Hardware (see "About hardware" on page 55) for more information.

Where can I see how many grace periods I have left? This information is available from the *License Information* page. When you add a new camera for which you have a license, you are granted a new full grace period for the camera in question from the date you added the camera. Therefore the end date of the grace period displayed on the *License information* page is for the first added but not activated camera.

Get additional licenses

What if you want to add - or if you already have added - more device channels (cameras or dedicated input/output boxes) than you currently have licenses for? In that case, you must buy additional licenses before the cameras will be able to send data to your system.

To get additional licenses for your system, contact your product vendor. In the short period until you get the additional licenses, you can disable some less important cameras (see "About hardware" on page 55) to allow some of the new cameras to run instead. When you have received an updated license file (*.lic*) with the new licenses, you must activate your licenses (see "Activate (Register) Licenses - Online or Offline" on page 44).

Licenses and camera replacement

You can replace a camera licensed in your system with a new camera, and have the new camera activated and licensed instead.

The total number of purchased device channels corresponds to the total number of cameras that are able to run on the surveillance system simultaneously. If you remove a camera from a recording server, you also free a license.

If you replace a camera with a similar camera (manufacturer, brand, and model), and give the new camera the same IP address as the old one, you will maintain full access to all the camera's databases. In this case, you move the network cable from the old camera to the new one without changing any settings in the Management Client, and then activate the license.

If replacing a camera (see "Replace hardware device" on page 56) with a different model, you must use the Management Client's *Replace Hardware* wizard to map all relevant databases of cameras, microphones, inputs, outputs, and so on. When done, remember to activate the license.

There is no limit to the number of cameras you can replace.

Licenses and OnSSI Federated Architecture?

Refer to OnSSI Federated Architecture Overview (see "About OnSSI Federated Architecture" on page 212).

Manage Software License Codes

When you purchase your system, you receive a Software License Code (SLC), which is used when installing your system.

Change Software License Code

Often your installation is run on a trial Software License Code (SLC) during the first period. When the trial period is over, and it is time to change the trial SLC to the permanent SLC, you can do this without any un- or reinstall action.

IMPORTANT: This must be done locally on the management server in question; you **cannot** do this from the Management Client.

1. On the management server, go to the notification area of the taskbar (a.k.a. *Systray*).

2. Right-click the *Management Server* icon, select *Change License...*.
3. The *Change License* dialog appears. Click *Import License...*.
4. Next, select the SLC license file saved for this purpose. When done, the selected license file location will be added just below the ***Import License...*** button.
5. Click *OK*. You are now ready to perform SLC registration.

Remote connect services

About remote connect services

Depending on the recording component, functionality described here may be limited or unavailable.

The remote connect services feature contains the Axis One-click Camera Connection technology developed by Axis Communications. It enables the system to retrieve video (and audio) from external cameras where firewalls and/or router network configuration normally prevents initiating connections to such cameras. The actual communication takes place via so-called secure tunnel servers (ST servers).

ST servers use a Virtual Private Network (VPN). Only devices holding a valid key can operate within a VPN. This offers a secure tunnel where data can be exchanged between public networks in a safe way.

Remote connect services allows you to:

- Edit credentials within the Axis Dispatch Service
- Add, edit, and remove ST servers
- Register/Unregister and edit Axis One-click cameras
- Go to the hardware related to the Axis One-Click camera.

Before you can use Axis One-click Camera Connection, you must first **install a suitable ST server environment**.

Install STS environment for One-click camera connection

1. Contact your system provider to obtain the needed user name and password for Axis Dispatch Services
2. Make sure your camera(s) support Axis Video Hosting System, http://www.axis.com/products/avhs/ (http://www.axis.com/products/avhs/).
3. If needed, update your Axis cameras with the newest firmware, http://www.axis.com/techsup/firmware.php (http://www.axis.com/techsup/firmware.php)
4. On each camera's homepage, go to *Basic Setup*, *TCP/IP*, and select *Enable AVHS* and *Always*
5. From your management server's download web page (see "Download Manager/download web page" on page 19) (controlled by the Download Manager), install the *Axis One-Click Connection Component* to setup a suitable Axis secure tunnel framework
6. From ***Services*** (search for *services.msc* on your machine), start the *Axis One-Click service* .

Edit Axis Dispatch Service properties

1. The Properties pane (see "Panes Overview" on page 33) displays relevant dispatch information on the *Axis Dispatch Service* tab.
2. Edit properties (see "Axis One-Click Camera connection properties" on page 51).
3. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

Add/edit STSs

1. Do one of the following:

a) To add an ST servers, right-click the *Axis Secure Tunnel Servers* top node, select *Add Axis Secure Tunnel Server...*

or

b) To edit an ST server, right-click it, select *Edit Axis Secure Tunnel Server...*

2. In the window that opens, fill in the relevant information (see "Axis One-Click Camera connection properties" on page 51).
3. If you chose to use credentials when you installed the *Axis One-Click Connection Component*, make sure to select the *Use credentials* check box and fill in exactly the same user name and password as used for the *Axis One-Click Connection Component*.
4. Click *OK*.

Remove STSs

1. To remove an ST server, right-click it, select *Remove Axis Secure Tunnel Server...*
2. Click *Yes*.

Register new Axis One-click camera

1. To register a camera under an ST server, right-click it, select *Register Axis One-click Camera...*
2. In the window that opens, fill in the relevant information (see "Axis One-Click Camera connection properties" on page 51).
3. Click *OK*.
4. The camera will now appear under the relevant ST server.

The color coding of the camera is either:

Color	Description
Red	Initial state—registered, but not connected to the ST server.
Yellow	Registered—connected to the ST server, but not added as hardware.
Green	Added as hardware—may or may not be connected to the ST server.

When added, status will always be green. The connection status (see "Read server service icons - management, recording and failover" on page 246) is then—as normal—reflected by *Devices* on *Recording Servers* in the Overview pane (see "Panes Overview" on page 33).

In the Overview pane (see "Panes Overview" on page 33), you may group your cameras for an easier overview.

If you choose **not** to register your camera at the Axis dispatch service at this point, you can do so later from the right-click menu (select *Edit Axis One-click Camera...*).

Unregister Axis One-click Camera

1. To unregister a camera under an ST server, right-click it, select *Unregister Axis One-click Camera.*
2. In the dialog that appears, make sure the check mark is selected and click *Yes*.
3. The camera will disappear from under the relevant ST server.

Axis One-Click Camera connection properties

Name	Description
Camera password	Enter/Edit. Provided with your camera at purchase. For further details, see your camera's manual or www.axis.com (http://www.axis.com).
Camera user	See details for *Camera password*.
Description	Enter/Edit a description of the item. Not compulsory.
External address	Enter/Edit the http address of the ST server where the camera(s) connect. **Tip:** Remember *http://* in front of the address.
Internal address	Enter/Edit the http address of the ST server where the recording server connects. **Tip:** Remember *http://* in front of the address.
Name	If needed, edit the name of the item.
Owner authentication key	See *Camera password.*
Passwords (for Dispatch Server)	Enter password. Must be identical to the one received from your system provider.
Passwords (for ST server)	Enter password. Must be identical to the one entered when the *Axis One-Click Connection Component* was installed.
Register/Unregister at the Axis Dispatch Service	Indicate whether you wish to register your Axis camera with the Axis dispatch service. Can be done at time of setup or later.
Serial number	(only relevant for hardware) Hardware serial number as specified by the manufacturer. The serial number is often, but not always, identical to the MAC address.
Use credentials	If it was decided—during installation of the ST server—to use credentials, select the check box.
User name (for Dispatch Server)	Enter user name. Must be identical to the one received from your system provider.
User name (for ST server)	Enter user name. Must be identical to the one entered when the *Axis One-Click Connection Component* was installed

Servers and hardware

Add hardware

The *Add Hardware* wizard helps you detect IP hardware devices, such as cameras and video encoders, on your network and add them to recording servers on your system.

1. To access *Add Hardware*, expand the *Servers* folder in the Management Client's Site Navigation pane (see "Panes Overview" on page 33) and select the *Recording Server* node.
2. In the Overview pane (see "Panes Overview" on page 33), right-click the required recording server and select *Add Hardware...*

The wizard offers you several ways of detecting and adding hardware devices:

Name	Description
Express (Recommended)	The system scans automatically for available hardware on the recording server's local network. **Tip:** If you are new to the system then use the *Express* hardware detection as it will guide you through each of the steps involved in detecting and adding your IP devices.
Address range scanning	The system scans defined network IP address ranges and detects hardware models.
Manual	Specify the IP address and port for each device. Cannot be used for adding remote systems in OnSSI Interconnect setups.
Remote connect hardware	Add hardware connected via a remotely connected server. All options offer the possibility of automatically detecting the correct hardware drivers Cannot be used for adding remote systems in OnSSI Interconnect setups.

It is strongly advised that you **only** add a physical hardware device to **one recording server** at a time.

Express

The *Express (recommended)* option automatically discovers hardware models on the recording server's local network.

1. Select *Express (recommended)* and click *Next*.
2. Specify user names and passwords if your hardware devices are not using the factory default user name and password. You can add as many user names and passwords as required by clicking *Add*. Remember to select the *Include* check box for each required device. When ready, click *Next*.
3. Wait while the hardware is detected. A status indicator will show the detection process. Once detection is complete, click *Next*.
4. Wait while device-specific information is collected for each hardware device. A status indicator shows the detection process. If collecting hardware information for a device is unsuccessful, click the *Failed* error message to see why. Once collection is complete, click *Next*.
5. Choose to enable or disable successfully detected hardware and cameras. Detected hardware, such as hardware device, camera, microphone and speaker is listed individually, allowing you to, for example, add a hardware device's camera without enabling its speaker if needed.
6. Select a default group for all device types, or group the devices individually. The devices are listed according to type, for example, camera, microphone, speaker. Click *Finish*.

Tip: Select the *Show hardware running on other recording servers* check box to see if detected hardware is running on other recording servers.

Address Range Scanning

The *Address Range Scanning* option scans your network for relevant hardware devices and OnSSI Interconnect remote systems (see "About OnSSI Interconnect" on page 57) based on your specifications regarding required IP ranges, drivers, and device user names and passwords.

1. Select *Address Range Scanning* and click *Next*.
2. Specify user names and passwords if your hardware devices are not using the factory default user name and password. You can add as many user names and passwords as required by clicking *Add*. Remember to select the *Include* check box for each required device. You must add and include at least one user name and password in order for the wizard to continue. When ready, click *Next*.
3. Select which drivers to use when you scan. By default, the system uses all known drivers. If your organization only uses certain hardware devices and/or models, you can achieve faster scanning by selecting only the drivers required for those hardware devices. Click *Next*.
4. Specify the IP address network ranges you want to scan for hardware.
   - **Start address**: First IP address in required range.
   - **End address:** Last IP address in required range. The start and end IP address may be identical, allowing you to only scan for a single hardware device if needed.
   - **Port:** Port number(s) on which to scan. Default is port 80. If your hardware devices are located behind a NAT-enabled router or a firewall, you may need to specify a different port number. When this is the case, also remember to configure the router/firewall so it maps the port and IP addresses used by the hardware devices.

   You can add as many network ranges as needed by clicking *Add* to add another row. You can add any network address between 0.0.0.1 and 255.255.255.255. You must select at least one network range before you can continue. Remember to select the *Include* check box for each required range.

   Wait while the hardware is detected. A status indicator shows the detection process. If you successfully detect hardware on a specified network range, a *Success* message appears in the *Status* column. If you fail to add a network range, you can click the *Failed* error message to see why. Once detection is complete, click *Next*.
5. Wait while device-specific information is collected for each hardware device. A status indicator shows the detection process. If collecting hardware information for a device is unsuccessful, click the *Failed* error message to see why the collection of information has failed. Once collection is complete, click *Next*.
6. Choose to enable or disable successfully detected hardware and cameras. Detected hardware, such as hardware device, camera, microphone and speaker is listed individually. This allows you to, for example, add a hardware device's camera without enabling its speaker if needed.
7. Select a default group for all device types. The devices are listed according to type, for example, camera, microphone and speaker. Click *Finish.*

Tip: The list of drivers that appears when you scan for drives is typically very long, and all drivers are selected by default. With *Select All* and *Clear All*, you can avoid having to select/clear all check boxes manually. Furthermore, when hardware is being detected, select the *Show hardware running on other recording servers* check box to see if detected hardware is running on other recording servers. Note also, that you can only specify IPv4 addresses when using *Address Range Scanning*.

Manual

The *Manual* option lets you specify details about each hardware device and OnSSI Interconnect remote systems (see "About OnSSI Interconnect" on page 57) separately. This can be a good choice if you only want to add a few hardware devices, and you know their IP addresses, required user names and passwords, and so on,

1. Select *Manual* and click *Next*.
2. Specify user names and passwords if your hardware devices are not using the factory default user name and password. You can add as many user names and passwords as required by clicking *Add*. Remember to

select the *Include* check box for each required device. You must choose to add and include at least one user name and password in order for the wizard to proceed. When ready, click *Next*.

3. Now select which drivers to use when scanning. By default, the system will use all known drivers. If your organization only uses certain hardware devices and/or models, you can achieve faster scanning by selecting only the drivers required for those hardware devices. When ready, click *Next*.
4. Specify information for the hardware you want to add. You can also optionally select the type of driver you want to add to speed up hardware detection.
   - ***Address:*** Specify the IP address of the hardware, you want to add.
   - ***Port:*** Specify the port number to which the camera is added.
   - ***Hardware driver:*** Select the driver of the hardware you want to add. Or select *Auto-detect* to let the wizard detect which driver to install.
5. Wait while the hardware is detected. A status indicator will show the detection process. Select or clear the network ranges to use in the detection process. If you successfully detect hardware, a *Success* message will appear in the *Status* column. If you fail to add a network range, click the *Failed* error message to see why. Once detection is complete, click *Next*.
6. Choose to enable or disable successfully detected hardware and cameras. Detected hardware, such as hardware device, camera, microphone and speaker is listed individually, allowing you to, for example, add a hardware device's camera without enabling its speaker if needed.
7. Select a default group for all device types or group the devices individually. The devices are listed according to type, for example, camera, microphone, speaker. Click *Finish.*

Tip: The list of drivers that appears when you scan is typically very long, and by default all drivers are selected. With *Select All* and *Clear All* , you can avoid having to select/clear all check boxes manually.

Tip: Select the *Show hardware running on other recording servers* check box to see if detected hardware is running on other recording servers.

Remote connect hardware

Remote Connect hardware automatically scans for hardware connected via a remotely connected server.

1. Select *Remote Connect hardware* and click *Next*.
2. Wait while the hardware is detected. A status indicator will show you how far you are in the detection process.
3. Once detection has completed, select which hardware you want to add and click *Next*.

About hardware

What is the difference between "hardware devices" and "hardware"? Technically, you do not add cameras or microphones to the system; rather you add **hardware devices**. This is because hardware devices have their own IP addresses or host names. Being IP-based, the system primarily identifies units based on their IP addresses or host names. Even though each hardware device has its own IP address or host name, several cameras, microphones, and so on can be attached to a single hardware device and share the same IP address or host name. This is typically the case with cameras attached to video encoder devices. You can of course configure and use each camera, microphone, and so on individually, even when several of them are attached to a single hardware device. **Hardware** on the other hand is a general term for cameras, microphones, and so on.

For each recording server on your system, you have several options for managing added IP hardware.

Most configuration and management of individual camera settings (see "Manage cameras" on page 82) (such as a camera's recording settings), input settings (see "Manage input" on page 109), and output settings (see "Manage output" on page 113) takes place on a more detailed level (camera, input or output level).

IMPORTANT: When you delete one or all hardware devices on a recording server, all its recordings are deleted permanently. If you need to add the hardware device to a recording server again, select the required recording server and use the Add Hardware (on page 53) wizard.

Edit basic hardware device settings

You are able to edit basic settings, such as IP address/host name, for added hardware:

1. In the Overview pane (see "Panes Overview" on page 33), expand the required recording server, right-click the hardware device you wish to edit.
2. From the menu that appears, select *Edit IP Hardware...* This opens the *Edit Hardware* window, where you can edit relevant properties (see "Specify hardware and device info properties" on page 119).
3. Click *OK*.

Replace hardware device

When you replace a hardware device on your network with another hardware device, you must know the IP address, port, user name and password of the new hardware device.

Furthermore, when replacing hardware devices, note that your system might be affected by license limitations (see "About licenses" on page 47). Using the ***Activate Online*** wizard (see "Activate (Register) Licenses - Online or Offline" on page 44), you must reactivate your licenses **after** replacing hardware devices. Also note, that if the new number of cameras, microphones, inputs and outputs exceeds the old number of cameras, microphones, inputs and outputs, you might also have to buy new licenses (see "About licenses" on page 47).

1. In the Overview pane (see "Panes Overview" on page 33), expand the required recording server, right-click the hardware device you wish to replace.
2. From the menu that appears, select *Replace Hardware*.
3. The *Replace Hardware* wizard appears. Click *Next*.
4. In the wizard, in the *Address* field (marked by red arrow in the image), enter the IP address of the new hardware. If known, select relevant hardware device driver from the *Hardware Driver* drop-down list (marked by red arrow in the image). Otherwise select *Auto Detect*. If port, user name or/and password data is different for the new device, also correct this **before starting the auto detect process (if needed)**.

   Tip: The wizard is prefilled with data from the existing hardware device. If you replace it with a similar hardware device, you can reuse some of this data - for example, port and driver information.
5. Do one of the following:
   - If you selected the required hardware device driver directly from the list, click *Next*.
   - If you selected *Auto Detect* in the list, click *Auto Detect*, wait for this process to be successful (marked by a ✓ to the far left), click *Next*.

   This step is designed to help you map devices and their databases, depending on the number of individual cameras, microphones, inputs, outputs and so on attached to the old hardware device and the new respectively.

   It is important to consider **how** to map databases from the old hardware device to databases of the new hardware device. You do the actual mapping of individual cameras, microphones, inputs, outputs, and so on by selecting a corresponding camera, microphone, input, output or *None* in the right-side column.

   > ***IMPORTANT:*** *Make sure to map **all** cameras, microphones, inputs, outputs, and so on. Contents stored in databases belonging to cameras, microphones, inputs, outputs, and so on mapped to* None*, are **lost**.*

   Click *Next*.
6. You are presented with a list of hardware to be added, replaced or removed. Click *Confirm*.

7. Final step is a summary of added, replaced and inherited devices and their settings. Click *Copy to Clipboard* to copy contents to an external source (for, for example, reporting purposes) or/and *Close* to end the wizard.

Disable/enable hardware device

Added hardware device is by default **enabled**.

In the Overview pane (see "Panes Overview" on page 33), under the required recording server, enabled/disabled hardware devices are indicated this way:

 Enabled

 Disabled

To disable added hardware device, for example, for licensing or performance purposes:

1. In the Overview pane, expand the required recording server, right-click the hardware device you wish to disable.
2. From the menu that appears, select *Enabled* to clear it.

Enable/disable individual devices

Cameras are enabled by default. **Microphones, speakers, inputs** and **outputs** are by default **disabled**.

This means that microphones, speakers, inputs and outputs must be individually enabled before they can be used on the system. The reason for this is that surveillance systems inherently rely on cameras, whereas the use of microphones and so on is highly individual depending on organizations' needs.

In the Overview pane (see "Panes Overview" on page 33), under the required server, enabled/disabled devices are indicated the following way (examples show indications for an output):

 Disabled

 Enabled

The same method for enabling/disabling is used for cameras, microphones, speakers, inputs, and outputs.

To enable a camera, input, or output:

1. In the Overview pane, expand the required recording server and the required hardware device. Right-click the camera, input, or output you wish to enable.
2. From the menu that appears, select *Enabled.*
3. Add the corresponding events to the camera, input, or output.

About OnSSI Interconnect

Depending on the recording component, functionality described here may be limited or unavailable.

OnSSI Interconnect allows you to integrate a number of smaller, physically fragmented, and remote installations with an Ocularis ES central site. These smaller sites, called remote sites, may even be installed on mobile units, for example, boats, busses or trains and may not even be permanently connected to a network.

Overview of OnSSI Interconnect

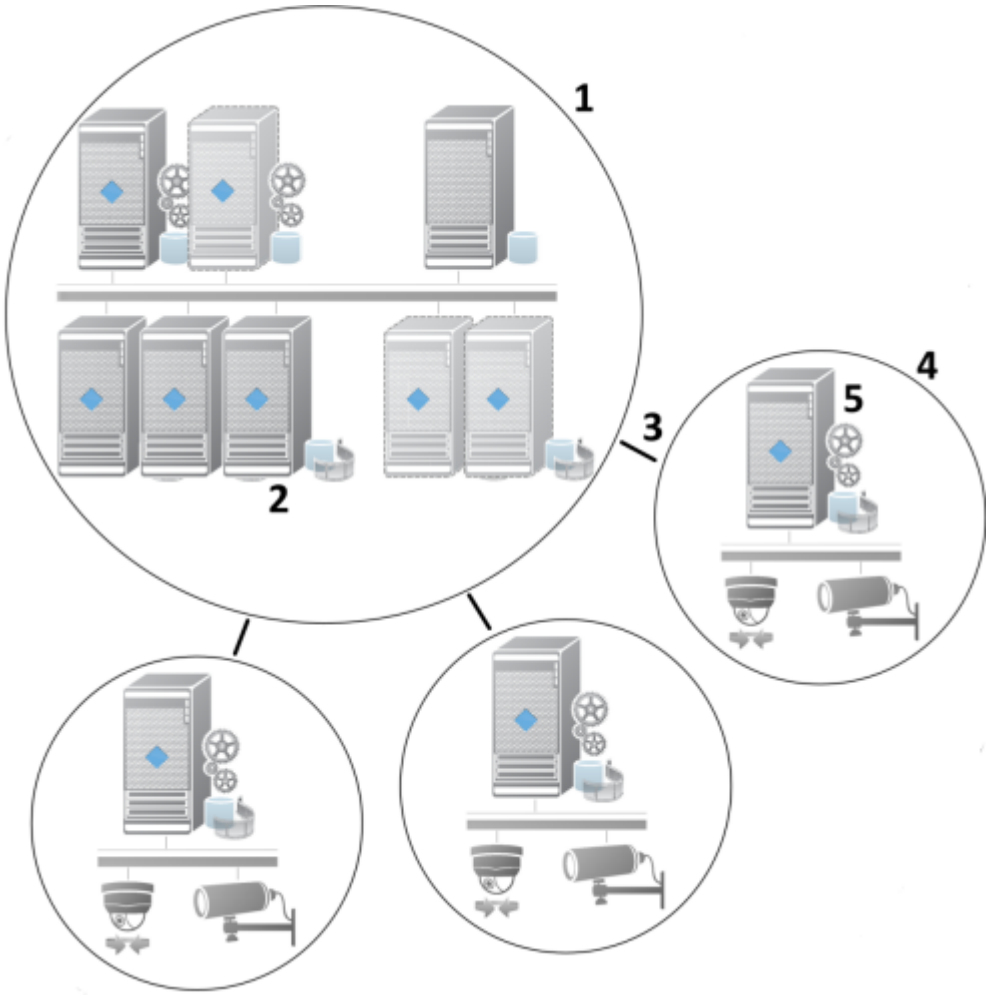

1. 1. OnSSI Interconnect central site
2. 2. OnSSI SMB* Ocularis VMS Products and Ocularis CS driver (handles the connection between the central sites' recording servers and the remote site, must be selected in the list of drivers when adding remote systems via the ***Add Hardware*** wizard)
3. 3. OnSSI Interconnect
4. 4. OnSSI Interconnect remote site (the complete remote site with system installation, users, cameras and so on)
5. 5. OnSSI Interconnect remote system (the actual technical installation at the remote site)

*Small and Medium-sized Businesses

More about OnSSI Interconnect

Each remote site runs independently and may perform any normal surveillance tasks. Depending on network connections and appropriate user rights, OnSSI Interconnect offers direct live viewing of remote site cameras and play back of remote site recordings from the central site. It also offers transfer of remote site recordings to the central site based on either events (see "Events overview" on page 161), rules/schedules (see "Manage rules" on page 165), or manual requests by Ocularis Client users. It also allows central site users to employ events originally triggered on remote sites (see "Events tab overview" on page 128) on the central site.

Which Ocularis product can act as central site and which can act as remote sites depends on the specific setup. Furthermore, it differs from setup to setup which versions, how many cameras, and how devices and events originating from the remote site are handled - if at all - by the central site.

Remote sites are added to the central site in the same way as multi-channel video encoders by use of the ***Add Hardware*** wizard (see "Add hardware" on page 53). However, remote sites can only be added using the ***Address range scanning*** (on page 54) or ***Manual*** (on page 54) options in the ***Add Hardware*** wizard. When adding the remote site, you must specify an account on the remote site. This account can be either a basic user, local Windows user, or domain user. It is possible to reuse an existing user or create a new one for usage with OnSSI Interconnect. However a new user must be created on the remote system before creating the OnSSI Interconnect setup. Depending on the user rights for the selected user on the remote site, the central site will get access to all cameras and functions or a sub-set of them.

Three possible OnSSI Interconnect setups

There are many possible ways to run OnSSI Interconnect. In the following, the three most likely scenarios are described. How to run your setup depends on your network connection, whether you request playback, and whether you retrieve remote recordings and to what degree.

What is remote recording? *Remote recording* (also known as edge recording) is both a physical camera supporting edge storage and a remote recording system in an OnSSI Interconnect setup. To minimize loss if a network breaks down, some physical cameras are able to store recordings on their own local storage. Either on request or automatically (depending on settings), recordings can be retrieved from remote storages to the surveillance system when the network is re-established. To save bandwidth it is possible to set up rules regarding when to retrieves recordings.
With **remote systems,** the principle is the same. However, recordings are **not** retrieved from remote cameras' edge storages, but from remote systems' recording servers.

Direct playback from remote sites on request (good network connections):
The most straight forward setup. The central site is continuously on-line with its remote sites which send remote recordings on request. Central site users play back remote recordings directly from the remote sites. This requires use of the ***Play back recordings from remote system*** option (see "Playback - remote system" on page 125).

Rule- or Ocularis Client-based retrieval of selected remote recording sequences from remote sites (periodically limited network connections):
Used when selected recording sequences (originating from remote sites) should be stored centrally to ensure independence from remote sites. Independence is crucial in case of network failure or network restrictions. Configuring retrieval of remote recordings when the network connection is optimal (i.e. not used for other priority data) can be done from the ***Remote Recordings*** tab (see "Remote Retrieval tab" on page 132). Alternatively, remote recordings retrieval can be started from the Ocularis Client when needed or a rule can be set up. In some scenarios, remote sites are on-line and in others, off-line most of the time. This is often industry specific. For some industries it is common for the central site to be permanently on-line with its remote sites (for example a retail HQ (central site) and a number of shops (remote sites)). For other industries, like transportation, the remote sites are mobile (for example, busses, trains, ships, and so on) and only able to establish network connection randomly. Should the network connection fail during a commenced remote recording retrieval, the job continues at next given opportunity. Note that if an automatic retrieval—or request for retrieval from the Ocularis Client—is received outside the time interval specified on the ***Remote Retrieval*** tab, it will be accepted, but not started until the selected time interval is reached. New remote recording retrieval jobs will queue and start when the allowed time interval is reached. Pending remote recording retrieval jobs can be viewed from the System Dashboard's Current Tasks (see "About current task" on page 195).

After connection failure, missing remote recordings are per default retrieved from remote sites:
Uses remote sites like a recording server uses the edge storage on a camera (see "Remote recording - camera/remote system" on page 128). Typically, remote sites are on-line with their central site, feeding it a live stream that the central site records. Should the network fail for some reason, the central site will miss out on recording sequences. However, once the network is re-established, the central site automatically retrieves remote recordings covering the down-period. This requires use of the ***Automatically retrieve remote recordings when connection is restored*** option (see "Remote recording - camera/remote system" on page 128).

Naturally, you can mix any of the above solutions to fit your organizations special needs.

OnSSI Interconnect and licensing

Cameras under remote sites in a OnSSI Interconnect setup are listed on the ***License Information*** page (see "License information" on page 47) of the central site. They are listed according to the same rules as other devices and are named just like "normal" devices but with OnSSI Interconnect in front - like this:

- *OnSSI Interconnect Device Channels*
- *OnSSI Interconnect Cameras*

Update remote site hardware

1. On the central site, in the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand ***Servers*** and select ***Recording Servers.***
2. In the Overview pane (see "Panes Overview" on page 33), expand the required recording server, select the relevant remote system. Right-click it.
3. From the menu that appears, select ***Update Hardware***. This opens the ***Update hardware*** dialog.
4. This dialog lists all changes (devices removed, updated and added) in the remote system since your OnSSI Interconnect setup was established or refreshed last. Click ***Confirm*** to update your central site with these changes.

Establish remote desktop connection to remote system

Preconditions: The remote desktop connections to the machine you want to remote to must be up and running and its management application must be open.

1. On the central site, in the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand ***Servers*** and select ***Recording Servers.***
2. In the Overview pane (see "Panes Overview" on page 33), expand the required recording server, select the relevant remote system.
3. In the Properties pane (see "Panes Overview" on page 33), select the ***Info*** tab.
4. In the ***Remote administration*** area, enter the appropriate Windows user name and password.
5. Once user name and password are saved, click ***Connect*** to establish remote desktop connection.
6. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

Enable playback directly from remote site camera

1. On the central site, in the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand ***Servers*** and select ***Recording Servers.***
2. In the Overview pane (see "Panes Overview" on page 33), expand the required recording server, select the relevant remote system. Select the relevant camera.
3. In the Properties pane (see "Panes Overview" on page 33), select the ***Record*** tab, and select the ***Play back recordings from remote system*** option (see "Playback - remote system" on page 125).
4. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

Note that in a OnSSI Interconnect setup, any privacy masking (see "Privacy Mask tab (camera properties)" on page 89) set on a remote system will be disregarded by the central system.

Retrieve remote recordings from remote site camera

1. On the central site, in the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand ***Servers*** and select ***Recording Servers.***

2. In the Overview pane (see "Panes Overview" on page 33), expand the required recording server, select the relevant remote system. Select the relevant camera.
3. In the Properties pane (see "Panes Overview" on page 33), select the ***Record*** tab, and select the ***Automatically retrieve remote recordings when connection is restored*** option (see "Remote recording - camera/remote system" on page 128).
4. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

As an alternative, you can use rules (see "Add a rule" on page 168) or start remote recording retrievals from the Ocularis Client when needed.

Note that in a OnSSI Interconnect setup, any privacy masking (see "Privacy Mask tab (camera properties)" on page 89) set on a remote system will be disregarded by the central system.

About storage and archiving

Depending on the recording component, functionality described here may be limited or unavailable.

When a camera or device records video and/or audio, all specified recordings are per default stored in the storage area defined for the device. More precisely in the storage area's default recording database named *Recording*. A storage area has no default archive(s), but these can easily be created.

Depending on recording settings, the storage area’s recording database will most likely run full at some point and its contents need to be archived in order to be saved. It is therefore possible to create archives within the default storage area and start an archiving process. Furthermore, it is possible to create alternative storage area(s) and configure that selected video/audio recordings must be stored/archived here.

Archiving is the automatic transfer of recordings from a camera's or device's default database to another location. This way, the amount of recordings you are able to store will not be limited by the size of the device's recording database. Archiving also makes it possible to back up your recordings on backup media of your choice.

Storage and archiving is configured on a per-recording server basis.

To ease explanations, the following mostly mentions cameras and video, but all is true about speakers and microphones and audio and sound as well.

IMPORTANT: We recommend that you use a dedicated hard disk drive for the recording server database. Using a dedicated hard disk drive for the database will prevent low disk performance. Furthermore, when formatting the hard disk, it is important to change its *Allocation unit size* setting from 4 to 64 kilobytes. This is to significantly improve recording performance of the hard disk. You can read more about allocating unit sizes and find help at http://support.microsoft.com/kb/140365/en-us.

IMPORTANT: The oldest data in a database will always be auto-archived (or deleted if no next archive is defined) when less than 5GB of space is free. If less than 1GB space is free, data will be deleted. A database always requires 250MB of free space; if this limit is reached (if data is not deleted fast enough), no more data will be written to the database until enough space has been freed. The actual maximum size of your database will thus be the amount of gigabytes you specify, minus 5GB.

Attaching devices to a recording server

Once you have configured the storage area and archiving settings for a recording server (where to store recordings, archives, how often to transfer recordings to archives, and so on), you can enable storage and archiving for individual cameras or a group of cameras (see "Attach a device or group of devices to storage area" on page 63). This is done from the individual devices or from the device group.

Effective archiving

When archiving is enabled for a camera or a group of cameras, the contents of the camera(s)' database will automatically be moved to an archive at regular intervals.

Depending on your requirements, you are able to configure one or more archives for each of your databases. Archives can be located either on the recording server computer itself, or at another location which can be reached by the system , for example on a network drive.

By setting up your archiving in an effective way, you can prune and groom your database storage area usage significantly if needed. Often, it is desired to make archived recordings take up as little space as possible—especially on a long-term basis, where it is perhaps even possible to slacken image and sound quality a bit. Effective pruning and grooming can help ensure this and can be handled from the *Storage* tab (see "Storage tab (recording server properties)" on page 69) of a recording server by adjusting several interdependent settings such as:

- Recording database retention
- Recording database size
- Archive retention
- Archive size
- Archive schedule
- Encryption
- Frames Per Second (FPS).

The size fields define the size of the camera's database, exemplified by the cylinder, and its archive(s) respectively:

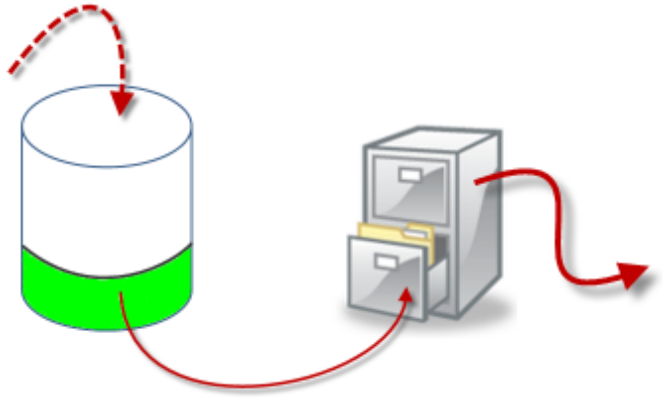

Recordings' way from recording database to archive to deletion

By means of retention time and size setting for the recording database, exemplified by the white area in the cylinder, you define how old recordings must be before they are archived. In our illustrated example, recordings are archived when they have "sifted" down into the green area of the database cylinder, or in other words: when they are old enough to be archived.

The retention time and size setting for archives define how long the recordings remain in the archive; recordings remain in the archive for the time specified, or until the archive has reached the specified size limit. When these settings are met, the system begins to overwrite old recordings in the archive.

The archiving schedule defines how often and at what times archiving takes place.

Encryption and FPS determine the size of the data in the databases.

To have recordings archived, all these parameters must be set up in accordance with each other. This means that the retention period of a next coming archive must always be longer than the retention period of a current archive or recording database. This is due to the fact that the number of retention days stated for an archive includes all retention stated earlier in the process. Furthermore, archiving must always take place more frequently than the retention period is set to, otherwise you risk losing data. If you have a retention time of 24 hours, any data older than 24 hours will be deleted. Therefore, to get your data safely moved to the next archive, it is important to run archiving more often than every 24 hours.

> ***Example:*** These storage areas (image to the left) have a retention time of 4 days and the following archive (image to the right) a retention time of 10 days. Furthermore, archiving is set to occur every day at 10:30, ensuring a much more frequent archiving than retention time.

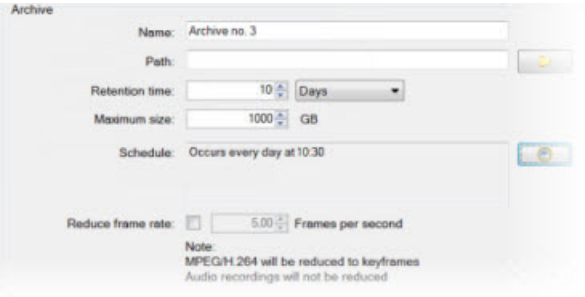

You can also control archiving by use of rules and events (see "About rules and events" on page 136).

Attach a device or group of devices to storage area

Once a storage area is configured for a recording server, you can enable it for individual devices (cameras, microphones or speakers) or a group of devices. You can also select which of a recording server's storage areas should be used for the individual device or the group.

1. In the Site Navigation pane (see "Panes Overview" on page 33), expand *Devices* and select either *Cameras*, *Microphones* or *Speakers* as required.
2. In the Overview pane (see "Panes Overview" on page 33), select the required device or a device group.
3. In the Properties pane (see "Panes Overview" on page 33), select the *Record* tab.
4. In the *Storage* area, select *Select...*
5. In the dialog that appears, select the wanted database, click *OK*.
6. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

View archived recordings

You view archived recordings in the Ocularis Client. As long as the archived recordings are stored locally or on accessible network drives, you can use the Ocularis Client's many features (timeline browser, evidence export, and so on) when browsing archived recordings; just like you would with recordings stored in a cameras' regular databases. The fact that you are viewing archived recordings are completely transparent.

Remember that individual user rights may prevent particular users from viewing recordings from particular cameras - just as is the case when browsing recordings from cameras' regular databases.

Back up archived recordings

Many organizations want to back up their recordings, using tape drives or similar. Exactly how you do this is highly individual, depending on the backup media used in your organization. However, the following is worth bearing in mind:

Back up archives rather than camera databases

Always create backups based on the content of archives, not based on individual camera databases. Creating backups based on the content of individual camera databases may cause sharing violations or other malfunctions.

When scheduling a backup, make sure the backup job does not overlap with your specified archiving times.

Tip: You are able to view each recording server's archiving schedule in each of a recording server's archives, on the Storage tab.

Knowing archive structure lets you target backups

When recordings are archived, they are stored in a certain sub-directory structure within the archive.

During all regular use of your system, the sub-directory structure will be completely transparent to the system's users, as they browse all recordings with the Ocularis Client regardless of whether the recordings are archived or not.

Knowing the sub-directory structure is primarily interesting if you want to back up your archived recordings (see "Archive structure" on page 64).

Archive structure

When recordings are archived, they are stored in a certain sub-directory structure within the archive.

During all regular use of your system, the sub-directory structure will be completely transparent to the system's users, as they browse all recordings with the Ocularis Client regardless of whether the recordings are archived or not. Knowing the sub-directory structure is primarily interesting if you want to back up your archived recordings.

In each of the recording server's archive directories, separate sub-directories are automatically created. These sub-directories are named after the name of the device and the name of the archive database.

Since you are able to store recordings from different cameras in the same archive, and since archiving for each camera is likely to be performed at regular intervals, further sub-directories are also automatically added.

These sub-directories each represent approximately an hour's worth of recordings. The one-hour split makes it possible to remove only relatively small parts of an archive's data if the maximum allowed size of the archive is reached.

The sub-directories are named after the device, followed by an indication of whether recordings come from an edge camera or via SMTP (if relevant), *plus* the date and time of the most recent database record contained in the sub-directory.

Naming structure:

```
...[Storage Path]\[Storage name]\[device-name] - plus date and time of most recent recording]\
```

If from edge camera:

```
...[Storage Path]\[Storage name]\[device-name] (Edge) - plus date and time of most recent recording]\
```

If from SMTP:

```
...[Storage Path]\[Storage name]\[device-name] (SMTP) - plus date and time of most recent recording]\
```

Real life example:

```
...F:\OurArchive\Archive1\Camera 1 on Axis Q7404 Video Server(10.100.50.137) - 2011-10-05T11:23:47+02:00\
```

Even further sub-directories are automatically added. The amount and nature of these sub-directories depend on the nature of the actual recordings. For example, several different such sub-directories will be added if the recordings are technically divided into sequences; something which is often the case if motion detection has been used to trigger recordings.

If you want to back up your archives, knowing the basics of the sub-directory structure enables you to target your backups.

Examples:

If wishing to back up the content of an entire archive, back up the required archive directory and all of its content; for example everything under:

```
...F:\OurArchive\
```

If wishing to only back up the recordings from a particular camera from a particular period of time, back up the contents of the relevant sub-directories only; for example everything under:

```
...F:\OurArchive\Archive1\Camera 1 on Axis Q7404 Video Server(10.100.50.137) - 2011-10-05T11:23:47+02:00\
```

Archive and virus scanning

If you are using virus scanning software on the computer on which the camera databases you want to archive are located, or on a computer to which data is archived, it is likely that the virus scanning will use a considerable amount of system resources on scanning all the data which is being archived.

This may affect system performance negatively. Also, virus scanning software may temporarily lock each file it scans, which may further impact system performance negatively.

If possible, you should therefore disable any virus scanning of camera databases and archiving locations.

Frequently asked questions about archiving

What happens if a storage area becomes unavailable? If a storage area becomes unavailable-for example if the storage area is located on a network drive, and the connection to the drive is lost— it will not be possible to store recordings in the storage area. Your system registers the availability of its recording servers' storage areas. This means that when a storage area becomes available again, it will again be possible to save recordings in the storage area. However, any recordings from the period in which the storage area was unavailable will be lost. When creating rules, you can use the events *Database Storage Area Unavailable* and *Database Storage Area Available* to trigger actions, such as the automatic sending of e-mail to relevant people in your organization. Furthermore, information about a storage area becoming unavailable/available will be logged.

How do I ensure that archiving is set up correctly? Archives are set up by adjusting several interdependent parameters correctly as described previously.

Can I create an archive on a network drive? Archives can be located either on the recording server computer itself, or at another location which can be reached by the system, for example on a network drive.

What happens when the maximum size of an archive is reached? When you create archives from the *Storage* tab, you specify a maximum size limit for the archive, in days and gigabytes. When either of the two maximum limits is reached, recordings in excess of the specified number of days/gigabytes will be removed. However, in order not to remove more recordings than necessary, excess recordings will be removed in chunks of approximately one hour's worth of recordings.

What happens if a scheduled archiving fails? If a scheduled archiving fails, for example because the archive is located on a network drive which is temporarily unavailable, the system will retry archiving after one minute. If that fails, another retry will take place after yet another minute, and so forth.
If the time of the next scheduled archiving is reached between two retries, an archiving attempt will be made at the scheduled time; if that attempt fails, the system will retry archiving after an hour, and so forth.

What happens if archiving is not finished before the next scheduled archiving? Your system inserts a compulsory period of archiving-free time after each finished archiving job. This ensures that archiving jobs do not overlap in time.

About recording servers

Recording servers are used for recording video feeds, and for communicating with cameras and other devices. A surveillance system will typically contain several recording servers, although only a single recording server is required for the system to work.

Recording servers on your system— i.e. computers with the recording server software installed, and configured to communicate with a management server— will be listed in the Management Client's Overview pane (see "Panes Overview" on page 33) when you expand the *Servers* folder in the Site Navigation pane (see "Panes Overview" on page 33) and then select the *Recording Servers* node.

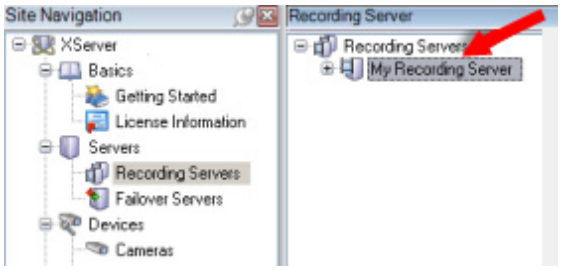

Recording server listed in Overview pane

Backward compatibility with recording servers from product versions older than this current version is limited. You can still access recordings on such older recording servers; but in order for you to be able to change their configuration, they must be of the same version as this current one. OnSSI highly recommends that all recording servers in your system are upgraded (see "Upgrade from previous version" on page 25) to the latest possible version.

IMPORTANT: When the **Recording Server service** is running, it is **very** important that neither Windows Explorer nor other programs are accessing Media Database files or folders associated with your system setup. Otherwise, the recording server might not be able to rename or move relevant media files. Unfortunately, this might bring the recording server to a halt. If this situation has already occurred, stop the Recording Server service, close the program accessing the media file(s) or folder(s) in question, and simply restart the Recording Server service.

Authorize a recording server

When first using the system, or when new recording servers have been added to the system, you must authorize the new recording servers.

Why must I authorize recording servers? In a surveillance system, recording servers point to management servers, not the other way round. In theory, recording servers which you do not want to include in your surveillance system could thus be configured to connect to your management servers. By authorizing recording servers before they can be used, surveillance system administrators have full control over which recording servers are able to send information to which management servers.

1. Expand the *Servers* folder in the Management Client's Site Navigation pane (see "Panes Overview" on page 33) and select the *Recording Servers* node.
2. Right-click the required recording server in the Overview pane (see "Panes Overview" on page 33).
3. From the menu that appears, select *Authorize Recording Server*:

After a short moment, the recording server will be authorized and ready for further configuration.

View/edit a recording server's properties

When a recording server is authorized, you are able to view/edit the recording server's properties, including its database storage area settings:

When you select the required recording server in the Management Client's Overview pane (see "Panes Overview" on page 33), the recording server's properties are displayed in the Properties pane (see "Panes Overview" on page 33). Expand the required recording server to see which devices are connected to the recording server. While the

Management Client loads information about the recording server, the text *(... expanding)* is displayed next to that recording server:

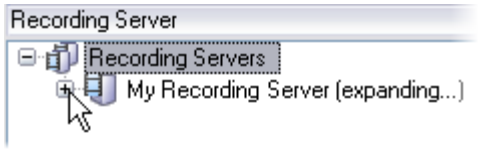

Add hardware to a recording server

You add IP hardware, such as cameras, video encoders, etc., to recording servers in your system through the *Add Hardware* wizard. The wizard helps you scan your network for relevant hardware. Refer to the wizard Add hardware (on page 53) for more information.

Manage hardware on a recording server

You have several options for managing hardware, such as cameras, video encoders, and so on, on recording servers in your system, refer to About devices (on page 82).

Remove a recording server

IMPORTANT: Removing a recording server will remove all configuration specified for the recording server through the Management Client, including all of the recording server's associated hardware (cameras, input devices, and so on).

1. Expand the *Servers* folder in the Management Client's Site Navigation pane (see "Panes Overview" on page 33) and select the *Recording Servers* node.
2. Right-click the no longer required recording server in the Overview pane (see "Panes Overview" on page 33).
3. From the menu that appears, select *Remove Recording Server*.
4. You will be asked to confirm that you want to remove the recording server and all of its associated hardware from the system. If you are sure, click *Yes*.
5. The recording server and all of its associated hardware will be removed.

Replace a recording server

If a recording server is malfunctioning and you want to replace it with a new server, while letting the new server inherit the settings of the old, malfunctioning recording server, do the following:

1. Retrieve the recording server ID from the old recording server:

a) In the Management Client's Site Navigation pane (see "Panes Overview" on page 33) select *Recording Servers*, then in the Overview pane (see "Panes Overview" on page 33) select the old, required recording server.

b) In the Management Client's Properties pane (see "Panes Overview" on page 33), select the *Storage* tab.

c) Press and hold down the CTRL key on your keyboard while selecting the *Info* tab.

d) Copy the recording server ID found in the lower part of the *Info* tab. Do not copy the term *ID* = but only the ID-number itself.

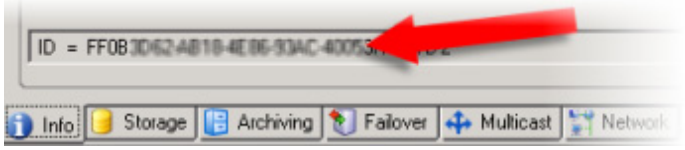

> ***IMPORTANT:*** *Stop the Recording Server service (see "Management Server service and Recording Server service" on page 244) on the old recording server, then in Windows'* Services *set the service's* Startup type *to* Disabled.

2. Replace recording server ID on the new recording server:

a) Make sure that the Recording Server service is stopped (see "Management Server service and Recording Server service" on page 244) and disabled on the old recording server.

It is very important that you do not start two recording servers with identical IDs at the same time.

b) On the new recording server, open an explorer and go to *C:\ProgramData\OnSSI* or the path where your recording server is located.

c) Open the file *RecorderConfig.xml*.

d) Delete the ID stated in between the tags *<id>* and *</id>*.

e) Paste the copied recording server ID in between the tags *<id>* and *</id>*. Save the *RecorderConfig.xml* file.

f) Restart the Recording Server service. When the new Recording Server service starts up, the recording server has inherited all settings on the old recording server.

Tip: This procedure also applies if you re-install Windows on the computer running the recording server, even if you do not replace the computer running the recording server.

Info tab (recording server properties)

You are able to verify or edit the name and description of a selected recording server on the *Info* tab. To access the *Info* tab, select the required recording server in the Overview pane (see "Panes Overview" on page 33), then select the *Info* tab in the Properties pane (see "Panes Overview" on page 33).

Info tab, displaying information about a recording server.

INFO TAB PROPERTIES

Name	Description
Name	Name of the recording server. The name will be used whenever the recording server is listed in the system and clients. A name is not compulsory, but highly recommended. The name does not have to be unique. To change the name, overwrite the existing name and click *Save* in the toolbar (see "Management Client Overview" on page 30). **Tip:** If you change the name, it will be updated throughout the system. This means that if the name is used in, for example, a rule, the name will automatically change in the rule as well.
Description	Description of the recording server. The description will appear in a number of listings within the system. For example, the description will appear when pausing the mouse pointer over the recording server's name in the Overview pane (see "Panes Overview" on page 33). A description is not compulsory. To specify a description, type the description and click ***Save*** in the toolbar (see "Management Client Overview" on page 30).
Host name	Non-editable field, displaying the recording server's host name.
Web server URL	Non-editable field, displaying the URL of the recording server's web server. The web server is used, for example, for handling PTZ camera control commands, and for handling browse and live requests from Ocularis Clients. The URL will include the port number used for web server communication (typically port 7563).
Time zone	Non-editable field, displaying the time zone in which the recording server is located.

Storage tab (recording server properties)

Depending on the recording component, functionality described here may be limited or unavailable.

On the *Storage* tab, you are able to setup, manage and view storage areas for selected recording servers. Refer to About storage and archiving (on page 61) for at more general introduction to recoding an archiving.

What is a storage area? A storage area is a directory in which database content— primarily recordings from the cameras connected to the recording server— is stored in at least a recording database and possibly archived in a number of archiving databases. A default storage area with a default recording database is automatically created for each recording server when the recording server is installed on the system. Unless you specifically define that another storage area should be used for particular cameras, recordings from connected cameras are stored in individual camera databases in the recording server's default storage area. Archives can be added to a storage area at any time convenient.

To access the *Storage* tab, select the required recording server in the Overview pane (see "Panes Overview" on page 33), then select the *Storage* tab in the Properties pan (see "Panes Overview" on page 33)e.

It is **not** possible to add databases or edit a storage area if the recording server is offline.

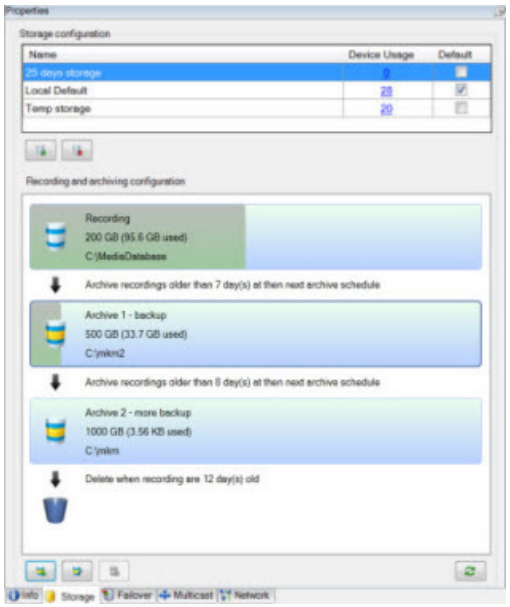

STORAGE TAB PROPERTIES

Storage configuration list contents: Name	Description
Name:	Indicates the name of the storage area. Click to edit.
Device Usage:	Indicates how many devises use the storage. Click the number link to see device details: 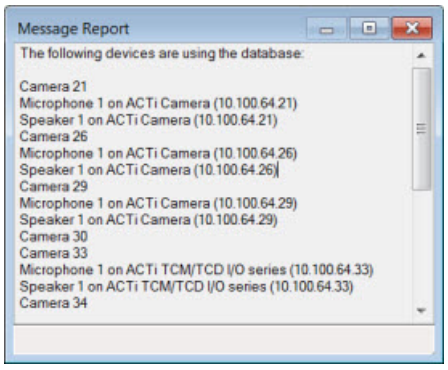
Default:	Indicates the default storage, that is the storage area in which database content is automatically stored unless you specifically define other storage areas for particular cameras. Only one storage at the time can be default.

Recording and archiving configuration list content:

1. Database name
2. Maximum size of the database (and usage; also represented graphically by a proportional filling of the database)

3. Database location
4. Archiving schedule for archiving to the next archive in the list.
   Note that the number of retention days stated for an archive includes all retention stated earlier in the process.

Tip: If you pause the mouse pointer over a database, it shows detailed database information.

ADD A STORAGE AREA

A storage area is always created with a predefined recording database named *Recording,* which you cannot rename. Apart from a recording database, a storage area can contain a number of archives (see "Create an archive within an existing storage area" on page 71).

1. To add an extra storage area to a selected recording server, click the  button located below the *Storage configuration* list.
2. This opens the *Storage and Recording Settings* dialog. Specify the relevant settings to continue:
3. Click *OK*.

If needed, you are now ready to create archive(s) within your new storage area (see "Create an archive within an existing storage area" on page 71).

CREATE AN ARCHIVE WITHIN AN EXISTING STORAGE AREA

A storage area has no default archive when it is created.

1. To create an archive, select the wanted storage area by clicking it in the *Recording and archiving configuration* list.
2. Next, click the  button located below the *Recording and archiving configuration* list.
3. This opens the *Archive Settings* dialog where you must specify the required settings (see "Storage and Recording settings" on page 71).

Click *OK*.

STORAGE AND RECORDING SETTINGS

In the Archive settings, specify the following:

Name	Description
Name	Rename the storage area if needed. Names must be unique.
Path	Type or use the browser link next to the field to specify the path to the directory in which to save the storage area. The storage area does not necessarily have to be located on the recording server computer itself. If the directory you plan to use does not already exist, you can create it using the browser dialog. Network drives must be specified using UNC (Universal Naming Convention) format, example: *server**volume**directory*\.

Name	Description
Retention time	Select a number of units and select either *Days* or *Hours* to specify how long recordings should stay in the archive before being deleted or archived (depending on archive settings). The retention time must always be longer than the retention time of the last archive or the recording database. This is due to the fact that the number of retention days stated for an archive includes all retention stated earlier in the process. **Example:** If you specify 24 hours, recordings must be at least a day old before they will be archived. If archiving is scheduled to take place before the 24 hours have passed, only recordings older than 24 hours will be archived. Bear in mind that the archive's scheduling may mean that recordings will be older than the specified number of hours before they are archived. This may especially be the case if you specify an archiving schedule with long time spans between archiving. Archiving is set up by adjusting several interdependent settings (see "About storage and archiving" on page 61).
Maximum size	Select the maximum number of gigabytes of recording data to save in the recording database. **Example:** If you want to store up to 100 gigabytes of recording data in the database, select 100. Recording data in excess of the specified number of gigabytes will be auto-moved to the first archive in the list - if any is specified - or deleted. **IMPORTANT**: This is one of two maximum size settings for the storage area. The *Retention Time* setting specified earlier may mean that recordings are removed from the archive before the specified number of gigabytes is reached. *IMPORTANT:* The oldest data in a database is always auto-archived (or deleted if no next archive is defined) when less than 5GB of space is free. If less than 1GB space is free, data is deleted. A database always requires 250MB of free space. If this limit is reached (if data is not deleted fast enough), no more data will be written to the database until enough space has been freed. The actual maximum size of your database will thus be the amount of gigabytes you specify, minus 5GB.
Schedule	Click the *Schedule* icon next to the *Schedule* field to specify an archiving schedule, that is the intervals with which the archiving process should start. If required, you can make archiving take place very frequently (in principle every hour all year round), or very infrequently (for example, every first Monday of every 36 months).
Reduce frame rate	Select the *Reduce frame rate* check box and set a frame per second (FPS) in order to reduce FPS when archiving. Reducing frame rates by a selected number of FPS's will make your recordings take up less space in the archive. On the other hand, it also reduces quality since a number of frames are erased, leaving only FPS corresponding to the number of FPS selected in the dialog. MPEG/H.264 will be reduced to minimum key-frames.

Tip: The ideal interval to use between each archiving process depends entirely upon your organization's needs. Consider your system's recording settings, make an estimate of the amount of data you expect to record within, for example, a day, a week, or a month, then decide on a suitable interval. Bear in mind that your organization's needs may change over time. It is a good idea to regularly monitor your archiving settings, and adjust them if required.

Tip: The effect of your selections is summed up in the lower part of the dialog. Use the summary to verify that your selections reflect your intentions.

Tip: If required, you can always adjust the archive's settings—including its scheduling—once the archive has been created.

Tip: It is possible to reduce frame rates to less than 1 FPS, for example as low as 0.1 FPS which means 1 frame every 10 seconds.

DELETE AN ARCHIVE FROM WITHIN AN EXISTING STORAGE

1. To delete an archive, select the wanted archive from the *Recording and archiving configuration* list by clicking it. A selected archive is marked by a dark frame.

   It is only possible to delete the last archive in the list. The archive does not have to empty.

2. Click the  button located below the *Recording and archiving configuration* list.
3. Click *Yes*.

DELETE AN ENTIRE STORAGE AREA

The storage area you want to delete must **not** be set as default storage area. Furthermore, it cannot be used by any devices to hold recordings. This means that you must possibly move devices and their **not yet archived** recordings to another storage area (see "Move non-archived recordings from one storage to another" on page 73) before you are allowed to delete the storage area.

1. Select the wanted storage area by clicking it.

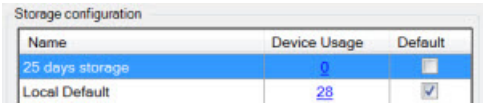

2. Click the  button located below the *Storage configuration* list.
3. Click *Yes*.

EDIT SETTINGS FOR A SELECTED STORAGE AREA OR ARCHIVE

1. In the *Recording and archiving configuration* list, to edit a storage area, select its recording database. To edit an archive, select the archive database.

   Tip: A selected database is marked by a dark frame.

2. Click the  button located below the *Recording and archiving configuration* list.
3. Either editing a recording database (see "Add a storage area" on page 71) or editing an archive (see "Create an archive within an existing storage area" on page 71).

If you change the maximum size of a database, recordings that exceed the new limit are auto-archived to the next archive or deleted - depending on archiving settings.

MOVE NON-ARCHIVED RECORDINGS FROM ONE STORAGE TO ANOTHER

1. Moving of contents from one recording database to another is done from the *Record* tab of the device in question.
2. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Devices,* select the wanted device type. In the Overview pane (see "Panes Overview" on page 33), select the wanted device.
3. In the Properties pane (see "Panes Overview" on page 33), click the *Record* tab. In the upper part of the *Storage* area, click *Select...* .
4. In the *Select Storage* dialog that follows, select the wanted database.
5. Click *OK*.

6. In the *Recordings Action* dialog that follows, select whether already existing - but **non-archived** -recordings should be moved along to the new storage or deleted.
7. After selecting, click *OK*.

   See also ***Record*** tab overview (on page 125).

ABOUT UPGRADING

Some information in this section may not be relevant due to differences in software versions.

If you are running Ocularis ES and your system is upgraded to Ocularis ES version 4.0 (or future versions), you might experience that you end up with a lot more storages than before upgrade. This is due to the fact that from version 4.0 and forwards, database structure is somewhat different than it used to be and during the update process, the system creates a number of extra databases. However, since your original naming-convention is respected, you can reconstruct your former database structure with only little moving about of devices and deletion of obsolete storages or databases.

Failover tab (recording server properties)

Depending on the recording component, functionality described here may be limited or unavailable.

If your organization uses failover recording servers, use the *Failover* tab to assign failover servers to recording servers. For any other details on failover recording servers, their settings, failover groups, and their settings, refer to About failover recording servers (see "About failover recording servers—regular and hot standby" on page 234).

ASSIGN FAILOVER RECORDING SERVERS

On the **Failover** tab of a recording server, you can choose between 3 different types of failover setups:

a No failover setup

b A primary/secondary failover setup

c A hot standby setup.

If you select **b** and **c**, you must select the specific server/groups. With **b**, you must also select a primary and optionally a secondary failover group. If the recording server becomes unavailable, a failover recording server from the primary failover group will take over. If you have also selected a secondary failover group, a failover recording server from the secondary group will take over in case all failover recording servers in the primary failover group are busy. This way you only risk not having a failover solution in the rare case when all failover recording servers in the primary, as well as in the secondary, failover group are busy.

1. In the Site Navigation pane (see "Panes Overview" on page 33), select *Servers*, *Recording Servers*. This opens a list of recording servers.
2. In the Overview pane (see "Panes Overview" on page 33), select the wanted recording server, go to the ***Failover*** tab.
3. To choose failover setup type (see "About failover recording servers—regular and hot standby" on page 234), select either ***None***, ***Primary failover server group***/***Secondary failover sever group*** or ***Hot standby server***. If relevant, select the needed server or groups from the dropdowns.

   You cannot select the same failover group as both primary and secondary failover group. Also regular failover servers already part of a failover group cannot be selected as hot standby servers.

   Tip: From the ***Primary***/***Secondary failover server group*** dropdowns, select ***Add new...*** to create new failover groups and add failover recording servers.
4. Next, click ***Advanced failover settings...***, this opens the ***Advanced Failover Settings*** window listing all devices attached to the selected recording server.

Tip: Even if you selected ***None***, ***Advanced failover settings*** will be available. Any selections are kept for later failover setups.

5. To specify the level of failover support, select ***Full Support***, ***Live Only*** or ***Disabled*** for each device in the list. Click ***OK***.
6. Finally, in the ***Failover service communication port (TCP)*** field, edit the port number if needed.

FAILOVER TAB PROPERTIES

- ***None:*** Select a setup without failover.
- ***Primary failover server group / Secondary failover sever group:*** Select a regular failover setup with one primary and possibly one secondary failover server group. Also, from the attached dropdown, select a primary failover group and possibly a secondary failover group.
- ***Hot standby server:*** Select a hot standby setup. Also, from the dropdown, select a hot standby server.
- ***Advanced failover settings...:*** Opens the ***Advanced Failover Settings*** window.
    - ***Full Support:*** Select to get full failover support for the device.
    - ***Live Only:*** Select to get live failover support for the device.
    - ***Disabled:*** Select to disable failover support for the device.
- ***Failover service communication port (TCP):*** By default, the port number is 11000. This port is used for communication between recording servers and failover recording servers. If changed, the recording server in question **must** be running and **must be** connected to the management server meanwhile.

Multicasting tab (recording server properties)

Your system supports multicasting of live streams from recording servers. In cases when many Ocularis Client users want to view live video from the same camera, multicasting can help save considerable system resources.

Multicasting is only possible for live streams; not for recorded video/audio.

If a recording server has more than one network interface card, it is only possible to multicast on one of them. Through the Management Client you are able to specify which one to use.

The successful implementation of multicasting also requires that your network equipment (switches, and so on) has been set up to relay multicast data packets to the required group of recipients only. If not; multicasting may not be different from broadcasting, which can significantly slow down network communication.

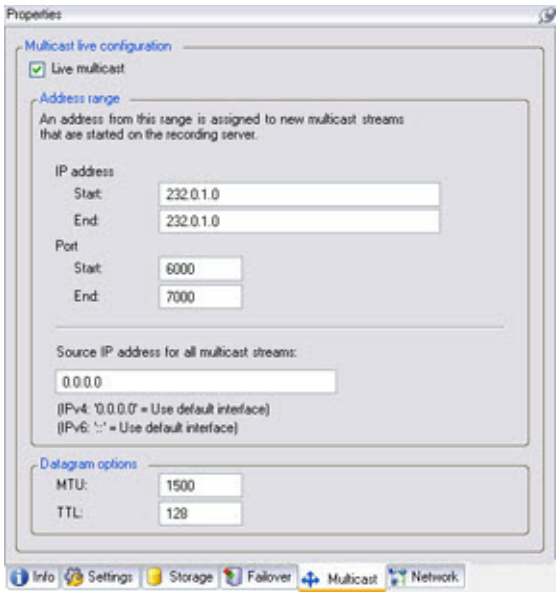

WHAT IS MULTICASTING?

In regular network communication, each data packet is sent from a single sender to a single recipient - a process known as unicasting. With multicasting, however, you can send a single data packet (from a server) to multiple recipients (clients) within a group. Multicasting can help save bandwidth.

- When using **unicasting**, the source must transmit one data stream for each recipient.
- When using **multicasting**, only a single data stream is required on each network segment.

Multicasting is an interesting option for streaming live video from recording servers to Ocularis Client s since video streams are not duplicated on each network segment.

Multicasting as described here is **not** streaming of video from camera to servers.

With multicasting, you work with a clearly defined group of recipients, based on options such as IP address ranges, the ability to enable/disable multicast for individual cameras, the ability to define largest acceptable data packet size (MTU), the maximum number of routers a data packet must be forwarded between (TTL), and so on. So, multicasting should not be confused with the much more primitive method *broadcasting*, which would send data to everyone connected to the network, even if the data is perhaps not relevant for everyone:

Name	Description
Unicasting	Sends data from a single source to a single recipient.
Multicasting	Sends data from a single source to multiple recipients within a clearly defined group.
Broadcasting	Sends data from a single source to everyone on a network; broadcasting can thus significantly slow down network communication.

WHAT ARE THE REQUIREMENTS?

To use multicasting, your network infrastructure must support IGMP (Internet Group Management Protocol, an IP multicasting standard). You must configure multicasting through the Management Client.

ENABLE MULTICASTING

On the *Multicast* tab, select the *Live multicast* check box. If the entire IP address range for multicast is already in use on one or more other recording servers, you cannot enable multicasting on further recording servers without freeing up some multicasting IP addresses first.

ASSIGN IP ADDRESS RANGE

In this section, you specify the range from which you want to assign addresses for multicast streams from the selected recording server. Clients connect to these addresses when viewing multicast video from the relevant recording server.

Name	Description
IP address	In the Start field, specify the first IP address in the required range. Then specify the last IP address in the range in the End field. For more info, see the following.
Port	In the Start field, specify the first port number in the required range. Then specify the last port number in the range in the End field.
Source IP address for all multicast streams	If a recording server has more than one network interface card, it is only possible to multicast on one of them. This field is therefore relevant if your recording server has more than one network interface card—or if it has a network interface card with more than one IP address. To use the recording server's default interface, leave the value 0.0.0.0 (IPv4) or :: (IPv6) in the field. If you want to use another network interface card, or a different IP address on the same network interface card, specify the IP address of the required interface.

SPECIFY DATAGRAM OPTIONS

In this section you specify settings for data packets (datagrams) transmitted through multicasting.

Name	Description
MTU	Maximum Transmission Unit, the largest allowed physical data packet size (measured in bytes). Messages larger than the specified MTU will be split into smaller packets before being sent. Default value is 1500, which is also the default on most Windows computers and Ethernet networks.
TTL	Time To Live, the largest allowed number of hops a data packet should be able to travel before it is discarded or returned. A hop is a point between two network devices, typically a router. Default value is 128.

ENABLE MULTICASTING FOR INDIVIDUAL CAMERAS

Even when you have specified multicasting settings for the selected recording server, multicasting will not work until you enable it for required cameras:

Select the required recording server in the Management Client's Site Navigation pane (see "Panes Overview" on page 33), select the required camera in the Overview pane (see "Panes Overview" on page 33), then select *Live multicast* on the Client tab (see "Client tab (camera properties)" on page 84) in the Properties pane (see "Panes Overview" on page 33). Repeat for all required cameras under the recording server in question.

SPECIFY IP ADDRESS RANGE

To specify the range from which you want to assign addresses for multicast streams from the selected recording server do the following:

For each multicast camera feed, the IP address/port combination (IPv4 example: 232.0.1.0:6000) must be unique. You can either use one IP address and many ports, or many IP addresses and fewer ports. By default, the system suggests a single IP address and a range of 1000 ports, but you can change this as required.

Example: If you want multicast for 1000 cameras, you would need either:

- 1 IP address and a range of 1000 different ports, OR
- a range of two IP addresses and a range of 500 different ports (or any matching combination), OR
- a range of 1000 IP addresses and a single port

When specifying the IP address, in the *Start* field, specify the first IP address in the required range. Then specify the last IP address in the range in the *End* field.

Tip: If required, a range may include only one IP address (IPv4 example: *232.0.1.0-232.0.1.0*)

IP addresses for multicasting must be within a special range set aside for dynamic host allocation by IANA (the authority overseeing global IP address allocation). If using IPv4, there is a certain range which goes from 232.0.1.0 to 232.255.255.255.

Network tab (recording server properties)

You define a recording server's public IP address on the *Network* tab. To access the *Network* tab, select the required recording server in the Overview pane (see "Panes Overview" on page 33), then select the *Network* tab in the Properties pane (see "Panes Overview" on page 33).

This description is also valid for failover recording servers (see "About failover recording servers—regular and hot standby" on page 234).

WHY USE A PUBLIC ADDRESS?

When an access client, such as an Ocularis Client, connects to a surveillance system, an amount of initial data communication, including the exchange of contact addresses goes on in the background. This happens automatically, and is completely transparent to users.

Clients may connect from the local network as well as from the internet, and in each case the surveillance system should be able to provide suitable addresses so the clients can get access to live and recorded video from the recording servers:

- When clients connect locally, the surveillance system should reply with local addresses (see "Manage local IP address ranges" on page 210) and port numbers.
- When clients connect from the internet, the surveillance system should reply with the recording server's public address, i.e. the address of the firewall or NAT (Network Address Translation) router, and often also a different port number (which is then forwarded to recording servers).

  To provide access to the surveillance system from outside a NAT (Network Address Translation) firewall, the system lets you use public addresses and port forwarding. This will allow clients from outside the firewall to connect to recording servers without using VPN (Virtual Private Network). Each recording server (and failover recording server) can be mapped to a specific port and the port can be forwarded through the firewall to the server's internal address.

ENABLE PUBLIC ACCESS

To enable public access, select the *Network* tab's *Enable public* access box.

DEFINE PUBLIC ADDRESS AND PORT

When public access is enabled, you can define the recording server's public address and public port number in the *Public address* and *Public port* fields respectively.

As public address, use the address of the firewall or NAT router which clients accessing the surveillance system from the internet must go through in order to reach recording servers.

Specifying a public port number is compulsory; it is always a good idea that port numbers used on the firewall or NAT router are different from the ones used locally.

When using public access, the firewall or NAT router used must be configured so requests sent to the public address and port are forwarded to the local address and port of relevant recording servers.

LOCAL IP RANGES

There are cases when the recording server's public address should not be used: When clients connect from the local network, the surveillance system should reply with local addresses and port numbers. The surveillance system must therefore be able to determine whether a client belongs on a local IP range or on the internet.

For this purpose, you are able to define a list of IP ranges (see "Manage local IP address ranges" on page 210) which the surveillance system should recognize as coming from a local network. On the *Network* tab, click *Configure...*.

Read recording server icons

The following icons are used in the Management Client to indicate the state of individual recording servers:

Recording Server service Icon	Description
	*Recording server is running*
	*Recording server is communicating*
	*Recording server requires attention* **:** This icon will typically appear because the Recording Server service has been stopped. **Tip:** You can verify whether the recording server is stopped by looking at the recording server icon in the notification area *of the computer running the recording server*. Right-clicking the recording server icon in the notification area opens a menu with which you can start/stop the Recording Server service, view recording server status messages, and so on. Refer to Recording Server service administration (see "Management Server service and Recording Server service" on page 244) for more information.
	*Recording server must be authorized* **:** Appears when the recording server is loaded for the first time. When first using a recording server, you must authorize it: Right-click the required recording server icon. From the menu that appears, select *Authorize Recording Server*. After a short moment, the recording server will be authorized and ready for further configuration.
	*Ongoing database repair* **:** Appears when databases have become corrupted, and the recording server is repairing them. The repair process may take considerable time if the databases are large. **IMPORTANT:** During the database repair it is not possible to record video from cameras connected to the recording server in question. Live video viewing will still be possible. **How can databases become corrupted?** Databases typically become corrupted if the recording server is shut down abruptly, for example due to a power failure or similar. Refer to Protect Databases from Corruption (see "Protect recording databases from corruption" on page 242) for useful information about how to avoid corrupt databases.

Change/verify a recording server's basic configuration

If you have installed several recording servers on your surveillance system, the recording servers should automatically be listed in the Management Client. If your Management Client does not list all the recording servers you have installed, the most likely reason is that the missing recording servers have not been correctly configured to

connect to a management server (in your system, recording servers point to management servers, not the other way round). The configuration normally takes place during one of the steps in the recording server installation process. Here, you specify recording server setup parameters, among these the IP address or host name of the management server to which the recording server should be connected. Fortunately, you do not have to re-install recording servers in order to specify which management servers they should connect to.

Once a recording server is installed, you can verify/change its basic configuration the following way:

1. On the computer running the recording server, right-click the ***Recording Server*** icon in the notification area:
2. From the menu that appears, select *Stop Recording Server service*:

   Important: Stopping the Recording Server service means that you cannot record and view live video while you verify/change the recording server's basic configuration.

3. Right-click the notification area's *Recording Server* icon again.
4. From the menu that appears, select *Change Settings...*:

   The *Recording Server Settings* window appears. Verify/change the following settings:

   - ***Management server hostname/IP address:*** Specify the IP address (example: *123.123.123.123*) or host name (example: *ourserver*) of the management server to which the recording server should be connected. This information is necessary in order for the recording server to be able to communicate with the management server.
   - ***Management server port:*** Specify the port number to be used when communicating with the management server. Default is port 9993, although you can change this if required.

5. Click *OK*.
6. To start the Recording Server service again, right-click the notification area's *Recording Server* icon, and select *Start Recording Server service*:

Tip: The notification area is occasionally also known as the system tray. It is located at the far right of the recording server computer's Windows taskbar.

Servers and clients require time-synchronization

Part of the security surrounding the use of clients with your system is based on time-based tokens.

Why servers require time-synchronization

When a client logs in to the surveillance system, the client receives a token from the management server. The token contains important security-related time information.

The management server also sends a similar token to the required recording server(s). This is partly due to the fact that recording servers may be located all around the world. Each recording server uses the token to validate the client's token against the local time in the recording server's own time zone.

The validity of a token expires after a while. Therefore, it is important that time on your management server and all of your organization's recording servers is synchronized (minute and second-wise; hours may of course be different in different locations around the world). If time on the servers is not synchronized, you may experience that a recording server is ahead of the management server's time.

When a recording server is ahead of the management server's time, it may result in a client's token expiring on the recording server earlier than intended by the management server. Under unfortunate circumstances, you might even experience that a recording server claims that a client's token has already expired when it receives it, effectively preventing the client from viewing recordings from the recording server.

How to synchronize time on your organization's servers depends on your network configuration, internet access, use of domain controllers, etc. Often, servers on a domain are already time-synchronized against the domain controller. If so, you should be fine as long as all required servers belong to the domain in question.

If your servers are not already time-synchronized, it will be necessary to synchronize the servers' time against a time server, preferably the same time server.

The following articles from Microsoft® describe what to do in different situations:

- How to configure an authoritative time server in Windows Server 2003
- Registry entries for the W32Time service

If these links do not work for you, try searching www.microsoft.com for time server, time service, synchronize servers or similar.

It is also very important that Ocularis Client s are time-synchronized with the management server.

Devices

About devices

You can either add (see "Add hardware" on page 53) or replace (see "About hardware" on page 55) devices.

In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), you are able to work with the following under *Devices*:

Name	Description
Cameras	Here you can handle the majority of camera configuration and management.
Microphones	On many devices you can attach external microphones. Some devices even have built-in microphones.
Speakers	On many devices you can attach external loudspeakers. Some devices even have built-in speakers.
Inputs	On many devices you can attach external units, typically external sensors, to input ports on the device. Input from such external input units can be used for many purposes in the system.
Outputs	On many devices you can attach external units to output ports on the device. This allows you to activate/deactivate lights, sirens, etc. through the system.

Manage cameras (on page 82)

Manage cameras

You can either add (see "Add hardware" on page 53) or replace (see "About hardware" on page 55) cameras.

Enabling/disabling as well as renaming of individual cameras takes place on the recording server hardware management level (see "About hardware" on page 55).

For all other configuration and management of cameras, expand *Devices* in the Management Client's Site Navigation pane (see "Panes Overview" on page 33), then select *Cameras*. In the Overview pane (see "Panes Overview" on page 33), you group your cameras for an easy overview of your cameras. Grouping also lets you specify common properties for all cameras within a group in one go and add cameras (see "About device groups" on page 104) to that group.

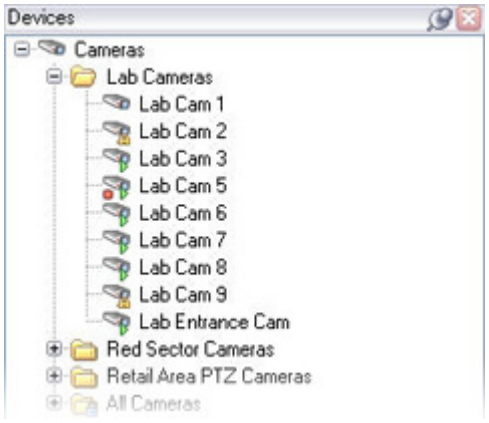

Device groups are used for grouping cameras

Once you have placed your cameras in groups, configuration can begin.

CONFIGURE INDIVIDUAL CAMERAS

You configure individual cameras by selecting the required camera in the list, then specifying the camera's required settings on the tabs in the Properties pane (see "Panes Overview" on page 33):

Tab	Use for specifying
Info (see "Info tab overview" on page 119)	The selected camera's name, etc
Settings (see "Settings tab overview" on page 121)	The selected camera's general settings.
Streams (see "Streams tab (camera properties)" on page 101)	The selected camera's video streams.
Record (see "Record tab overview" on page 125)	The selected camera's recording, database and archiving storage settings.
Presets (see "PTZ tab (video encoders)" on page 131)	The selected camera's preset positions (only available if the selected camera is a PTZ camera).
Patrolling (see "PTZ Patrolling tab (camera properties)" on page 93)	The selected camera's patrolling profiles (only available if the selected camera is a PTZ camera).
Events (see "Events tab overview" on page 128)	Events.
Client (see "Client tab (camera properties)" on page 84)	Information which will affect client's use of the selected camera.
Privacy Mask (see "Privacy Mask tab (camera properties)" on page 89)	Privacy masking for the selected camera.
Motion (see "Motion tab (camera properties)" on page 85)	The selected camera's motion detection settings.

READ THE CAMERA LIST'S STATUS ICONS

The following icons are used to indicate status of cameras (see "Manage cameras" on page 82), microphones (see "Manage Microphones" on page 102), speakers (see "Manage speakers" on page 107), input (see "Manage input" on page 109) and output (see "Manage output" on page 113) events in item lists:

Cam-era	Micro phone	Spea-ker	In-put	Out-put	Description
					**Item enabled:** Can communicate with the recording server, and can if required be started/stopped automatically through a rule.

Cam-era	Micro phone	Spea-ker	In-put	Out-put	Description
					**Item recording.** **Speaker being recorded:** Note that what is being said through the speaker can be recorded, but cannot subsequently be played back or exported (for example to prove that a warning was given).
					**Item temporarily stopped or has no feed:** Often shown when an item is communicating with the system while it is being disabled or enabled. Also shown if the Default Start Audio Feed Rule is not active; see Managing Rules. When stopped, no information is transferred to the system. In which case—if it is a camera—neither live viewing nor recording will be possible. However, a stopped item will still be able to communicate with the recording server, as opposed to when an item is disabled.
					**Item disabled:** Cannot be started automatically through a rule and will not be able to communicate with the recording server. In the case of a camera, when a camera is disabled, neither live viewing nor recording will be possible.
					**Item database being repaired.**
					**Item requires attention.**
					**Status unknown.**
					Note that some icons may be combined, as in this example where **Item is enabled** is combined with **Item is recording** (since a recording item is also an enabled item).

CLIENT TAB (CAMERA PROPERTIES)

The *Client* tab lets you specify information which will affect clients' use of the selected camera. To access the *Client* tab, select the required camera in the Overview pane, then select the *Client* tab in the Properties pane.

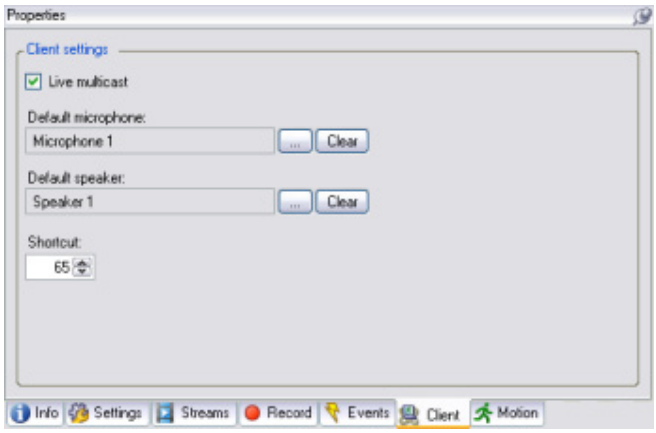

Client settings

Name	Description
Live multicast	The system supports multicasting (see "Multicasting tab (recording server properties)" on page 75) (sending of single data packets to multiple recipients within a group, thereby saving bandwidth and system resources) of live streams from recording servers to Ocularis Clients. To enable multicasting of live streams from the selected camera, select the check box. Remember that for the feature to work, multicasting (see "Multicasting tab (recording server properties)" on page 75) must also be configured for the recording server. If multicasting is not possible, for example due to restrictions on the network or on individual clients, the system will revert to unicasting (sending of separate data packets to separates recipients).
Default microphone	By defining a default microphone, you can determine from which microphone Ocularis Client users should by default hear recordings when they select the camera in question in their Ocularis Clients. The users can subsequently select another microphone if they require so. Bear in mind that although you have defined a default microphone for a camera, it cannot be guaranteed that all Ocularis Client users will hear audio from the microphone in question: Some users may not have speakers attached, some users may not have the rights required to listen to audio, etc.
Default Speaker	By defining a default speaker, you can determine through which microphone Ocularis Client users should by default be able to speak when they select the camera in question in their Ocularis Clients. The users can subsequently select another speaker if they require so. Bear in mind that although you have defined a default speaker for a camera, it cannot be guaranteed that all Ocularis Client users will be able to talk through the speaker in question: Some users may not have a microphone attached, some users may not have the rights required to talk through speakers, etc.
Shortcut	This field is not in use.

MOTION TAB (CAMERA PROPERTIES)

The *Motion* tab lets you enable and configure motion detection for the selected camera. Motion detection configuration is a key element in your system: Your motion detection configuration may determine when video is recorded, when events are generated, when external output (such as lights or sirens) is triggered, etc.

Time spent on finding the best possible motion detection configuration for each camera may therefore help you later avoid unnecessary alerts, etc. Depending on the physical location of the camera, it may be a very good idea to test motion detection settings under different physical conditions (day/night, windy/calm weather, etc.).

Before you configure motion detection for a camera, it is highly recommended that you have configured the camera's image quality settings, such as resolution, compression, etc., on the *Settings* tab (see "Settings tab overview" on page 121). If you later change image quality settings, you should always test any motion detection configuration afterwards.

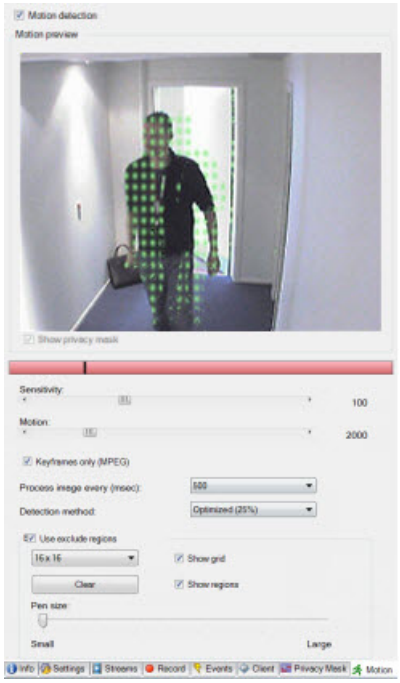

Camera properties: *Motion* tab with red deflection on the motion indication bar

You can configure motion detection for all cameras in a device group (see "Manage cameras" on page 82) in one go. Note, however, that some motion detection settings must be configured individually for each camera. This is the case with exclude regions (areas in which not to use motion detection), as these are very likely to vary from camera to camera.

Enable and disable motion detection

Motion detection is enabled by default. To enable/disable motion detection for a camera, select/clear the *Motion* tab's *Motion detection* check box.

When motion detection is disabled for a camera, any motion detection-related rules (see "Manage rules" on page 165) for the camera will not work.

Motion detection settings

You are able to specify settings relating to the amount of change required in a camera's video in order for the change to be regarded as motion. You are also able to specify intervals between motion detection analysis, any areas of an image in which motion should be ignored, etc.

Sensitivity slider

Determines **how much each pixel** in the camera's images must change before it is regarded as motion.

Drag the slider to the left for a higher sensitivity level, and to the right for a lower sensitivity level.

The *higher* the sensitivity level, the less change will be allowed in each pixel before it is regarded as motion.

The *lower* the sensitivity level, the more change in each pixel will be allowed before it is regarded as motion. This way you are able to allow insignificant changes, which should not be regarded as motion.

Pixels in which motion is detected are highlighted in green in the preview image. Select a slider position in which only detections you consider motion are highlighted.

Highlighted motion in the preview image

Tip: Your exact sensitivity slider setting is indicated by a number from 0-300 in the right side of the slider. This way you are able to compare the exact sensitivity slider setting between cameras.

Tip: If you find the concept of motion detection sensitivity difficult to grasp, try dragging the slider to the left towards the highest possible sensitivity (0) position: The more you drag the slider towards the highest possible sensitivity position, the more of the preview image becomes highlighted in green. This is because with a very high sensitivity level even the slightest change in each pixel will be regarded as motion.

Motion slider

Determines **how many pixels** in the camera's images image must change before it is regarded as motion.

The selected motion level is indicated by the black vertical line in the motion indication bar above the sliders.

The black vertical line in the motion indication serves as a threshold: When detected motion is above the selected sensitivity level, the bar changes color from green to red, indicating a positive detection.

Motion indication bar deflection changes color from green to red when above the threshold, indicating a positive motion detection

Tip: Your exact motion slider setting is indicated by a number from 0-10.000 in the right side of the slider. This way you are able to compare the exact motion slider setting between cameras.

Keyframe settings

Determines if motion detection should be done on keyframes only or on the entire video stream.

Select ***Keyframes only*** to do motion detection on keyframes only.

Image processing interval

Lets you select how often motion detection analysis should be carried out on video from the camera.

From the *Process image every (msec):* list, select the required interval: every 100 milliseconds (i.e. once every tenth of a second), every 250 milliseconds, every 500 milliseconds, every 750 milliseconds, or every 1000 milliseconds (i.e. once every second). Default is every 500 milliseconds.

The interval is applied regardless of the camera's frame rate settings.

Detection method

Lets you optimize motion detection performance by analyzing only a selected percentage of the image, for example 25%. By analyzing 25%, only every fourth pixel in the image is analyzed instead of all pixels.

Using optimized detection will reduce the amount of processing power used to carry out the analysis, but will also mean a less accurate motion detection.

In the *Detection method* drop down-box, select the wanted detection method.

Exclude regions

Lets you disable motion detection in specific areas of a camera's images. Parts of images in which motion should be ignored this way are called *exclude regions*.

Disabling motion detection in specific areas may help you avoid detection of irrelevant motion, for example if the camera covers an area where a tree is swaying in the wind or where cars regularly pass by in the background.

When exclude regions are used with PTZ cameras and you pan/tilt/zoom the camera, the excluded area will not move accordingly. This might mean that objects originally excluded will be included. This is due to the fact that the exclude region is locked according to the camera's view, not the excluded region.
Consequently, it is not recommended to use exclude regions with PTZ cameras.

To use exclude regions, select the *Use exclude regions* check box.

When done, the preview image will be divided into selectable sections by a grid.

To define exclude regions, drag the mouse pointer over the required areas in the preview image. Pressing down the left mouse button selects a grid section; right mouse button clears a grid section.

You are able to define as many exclude regions as you require. Excluded regions are shown in blue.

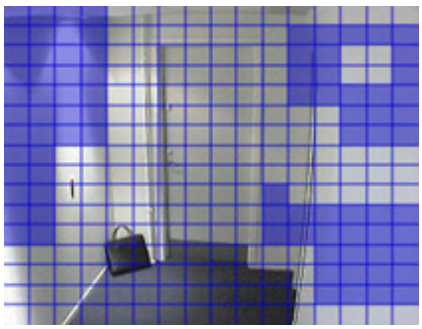

Three exclude regions defined in the preview window. In this case, the grid is visible.

The blue exclude area indications will only appear in the preview image on the *Motion* tab, not in any other preview images in the Management Client or access clients.

- ***Grid Size***
  The value selected in the *Grid size* list determines the density of the grid, regardless whether the grid is shown or not.

  Select between the values 8×8 (i.e. a grid dividing the image into eight sections along the X-axis and eight sections along the Y-axis), 16×16, 32×32 or 64×64.

  With a grid of 8×8, the image will be divided into relatively few sections for you to select for exclude regions. Each section will be relatively large; you will not be able to define very detailed exclude regions. With a grid size of 64×64, the image will be divided into relatively many sections for you to select for exclude regions. Each section will be relatively small, enabling you to define more detailed exclude regions.

  Examples of 8×8, 16×16, 32×32 and 64×64 grids respectively:

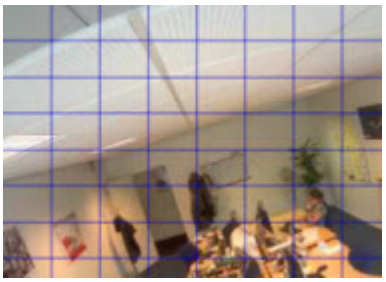

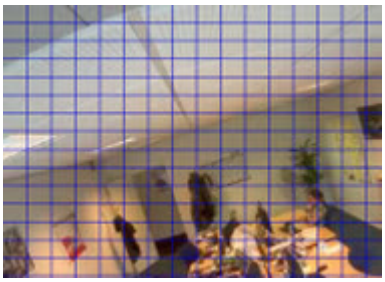

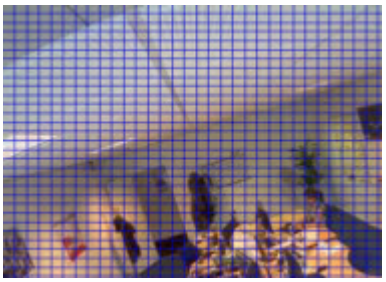

The four different grid sizes.

- ***Show Grid***

The grid may be visible or hidden, depending on whether the *Show grid* check box is selected or not.

When the *Show grid* check box is selected (default), the preview image will feature a grid indicating the division of the preview image into selectable sections. The grid may help you when selecting exclude regions in the preview image.

The density of the grid is determined by the value selected in the *Grid size* list.

Showing the grid is not a requirement for selecting exclude regions; even without the grid you are able to select exclude regions as described earlier. Hiding the grid may provide a less obscured view of the preview image.

- ***Show Regions***

When the *Show regions* check box is selected (default), exclude regions will be highlighted in blue in the preview image.

Hiding exclude regions may provide a less obscured view of the preview image. However, under normal circumstances it is highly recommended that you keep the *Show regions* box selected; otherwise exclude regions may exist without you or your colleagues being aware of it.

The blue exclude area indications will only appear in the preview image on the *Motion* tab, not in any other preview images in the Management Client or access clients.

- ***Pen size***

Use the *Pen size* slider to indicate the size of the selections you wish to make when clicking and dragging the grid to select regions for privacy masking. Default is set to small, which is equivalent to one square in the grid.

PRIVACY MASK TAB (CAMERA PROPERTIES)

The *Privacy Mask* tab lets you enable and configure privacy masking for the selected camera. Among other things, you can define if and how selected areas of a camera's view should be masked before distribution. For example, if a surveillance camera films a street, in order to protect residents privacy, you can mask certain areas of a building (could be windows and doors) with privacy masking. This is even needed in some countries to comply with national legislation.

As administrator you are also able to see through privacy masked areas, and can turn showing of privacy masked areas on and off. When viewed via Ocularis Client or any other media, privacy masked areas will be represented as black areas and it is impossible to see behind the privacy masking or in any way remove it.

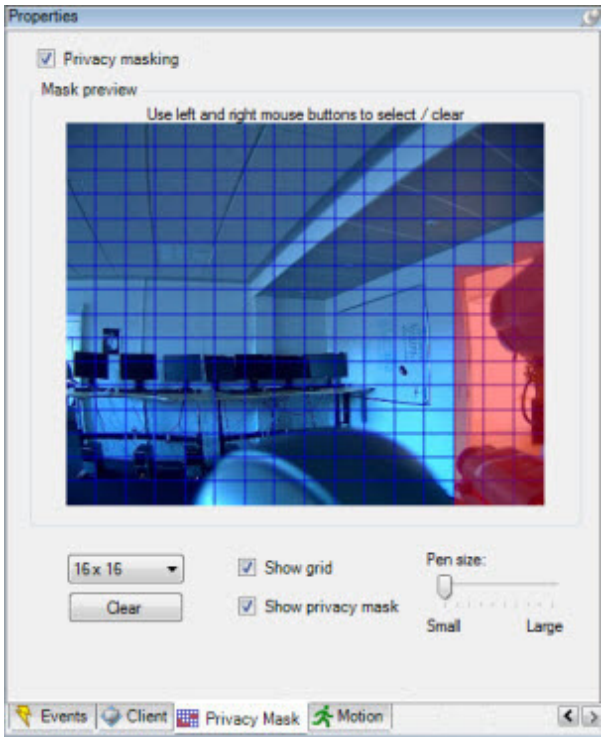

Red areas indicate the areas masked for privacy.

When privacy masks are used with PTZ cameras and you pan/tilt/zoom the camera, the selected area masked for privacy will **not** move accordingly. This might mean that objects masked for privacy become visible. This is due to the fact that the masked area is locked according to the camera's view, not the masked object. Consequently, it is not recommended to use privacy masking with PTZ cameras.

Enable and disable privacy masking

The privacy masking feature is enabled by default. To enable/disable the privacy masking feature for a camera, select/clear the *Privacy Mask* tab's *Privacy masking* check box.

Privacy masking settings

When you enable privacy masking, the preview image is divided into selectable sections by a grid.

To define privacy mask regions, drag the mouse pointer over the required areas in the preview image. Pressing down left mouse button selects a grid section; right mouse button clears a grid section.

You are able to define as many privacy mask regions as you require. Privacy mask regions are shown in red.

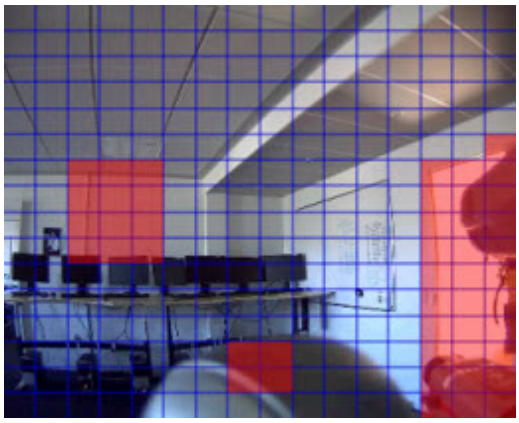

Three privacy mask regions defined in the preview window. In this case, the grid is visible.

The red privacy mask indications will also appear in the preview image on the *Motion* tab.

Name	Description
Grid Size	The value selected in the *Grid size* list determines the density of the grid, regardless whether the grid is shown or not. Select between the values 8×8 (i.e. a grid dividing the image into eight sections along the X-axis and eight sections along the Y-axis), 16×16, 32×32 or 64×64. With a grid of 8×8, the image will be divided into relatively few sections for you to select for privacy mask regions. Each section will be relatively large; you will not be able to define very detailed privacy mask regions. With a grid size of 64×64, the image will be divided into relatively many sections for you to select for privacy mask regions. Each section will be relatively small, enabling you to define more detailed privacy mask regions. Examples of 8×8, 16×16, 32×32 and 64×64 grids respectively:

	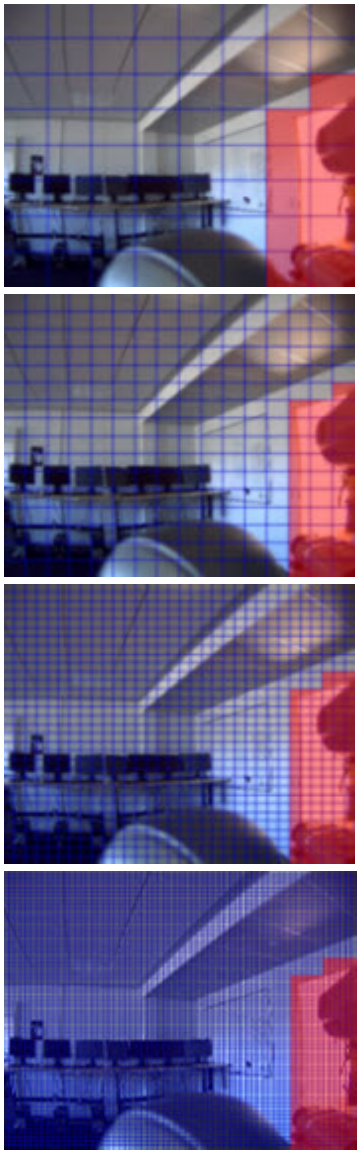 The four different grid sizes
Show Grid	The grid may be visible or hidden, depending on whether the *Show grid* check box is selected or not. When the *Show grid* check box is selected (default), the preview image will feature a grid indicating the division of the preview image into selectable sections. The grid may help you when selecting privacy mask regions in the preview image. Showing the grid is not a requirement for selecting privacy mask regions; even without the grid you are able to select privacy mask regions as described above. Hiding the grid may provide a less obscured view of the preview image.

Show Privacy Masks	When the *Show privacy masks* check box is selected (default), privacy mask regions will be highlighted in red in the preview image. Hiding privacy mask regions may provide a less obscured view of the preview image. However, under normal circumstances it is highly recommended that you keep the *Show privacy masks* box selected; otherwise exclude privacy mask regions may exist without you or your colleagues being aware of it.
Pen size	Use the *Pen size* slider to indicate the size of the selections you wish to make when clicking and dragging the grid to select regions for privacy masking. Default is set to small, which is equivalent to one square in the grid.

Privacy masking in OnSSI Interconnect

Note that in a OnSSI Interconnect setup, any privacy masking (see "Privacy Mask tab (camera properties)" on page 89) set on a remote system will be disregarded by the central system.

PTZ PATROLLING TAB (CAMERA PROPERTIES)

The *Patrolling* tab lets you create patrolling profiles, the automatic movement of a PTZ (Pan/Tilt/Zoom) camera between a number of preset positions (see "PTZ Presets tab (camera properties)" on page 96). Before you are able to work with patrolling, you must have specified at least two preset positions for the camera.

You manage patrolling on the *Patrolling* tab, which is available only when the selected camera is a PTZ camera.

Patrolling profiles are the definitions of how patrolling should take place. This includes the order in which the camera should move between preset positions, how long it should remain at each position for, etc. You are able to create an unlimited number of such patrolling profiles and use them in your rules (see "Manage rules" on page 165). For example, you may create a rule specifying that one patrolling profile should be used during daytime opening hours, and another during nights.

In order to use PTZ cameras' features, including the ability to pan, tilt, and zoom, operators must have a role which gives them the necessary rights. Refer to About roles (on page 182) for more information, including step-by-step descriptions of how to assign users to roles and how to specify the rights of roles.

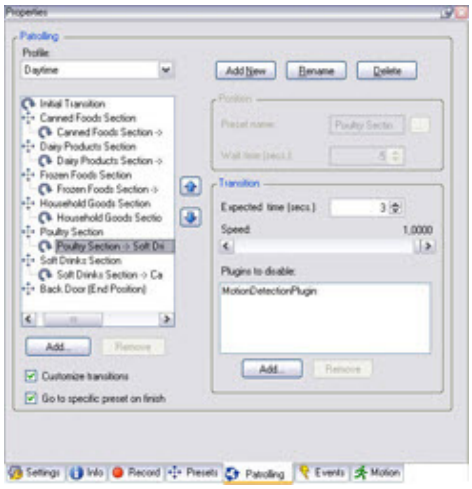

Patrolling tab, displaying a patrolling profile with customized transitions

Add a patrolling profile

1. Click *New*. This will open the *Add Profile* dialog.
2. In the *Add Profile* dialog, specify a name for the patrolling profile:

   Tip: Use a descriptive name; the name may later be used in situations where you will not have access to details about the item, e.g. when using it in a rule.
3. Click *OK*. The new patrolling profile will be added to the *Patrolling* tab's *Profile* list. You are now able to specify required preset positions and other settings for the patrolling profile.

Specify preset positions for use in a patrolling profile

1. Select the required patrolling profile in the *Profile* list:

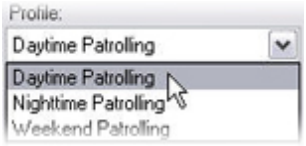

2. Click *Add* . This will open the *Select Preset* dialog.
3. In the *Select Preset* dialog, select the preset positions required for your patrolling profile:

4. Click *OK*. The selected preset positions are added to the list of preset positions for the patrolling profile:

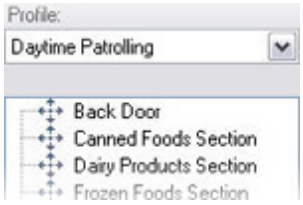

5. The preset position at the top of the list will be used as the first stop when the camera patrols according to the patrolling profile, the preset position in second position from the top will be the second stop, and so forth.

   If required, change the sequence by selecting the required preset position and using the up/down buttons:

   Tip: If required, you can easily add more preset positions to the list by clicking *Add*, or remove unwanted preset positions from the list by selecting the unwanted preset position, then clicking *Remove*.

Specify for how long to stay at each preset position

When patrolling, the PTZ camera will by default remain for 5 seconds at each preset position specified in the patrolling profile before it moves on to the next preset position. To change the number of seconds for which the PTZ camera will remain at a specific preset position, do the following:

1. Select the required patrolling profile in the *Profile* list.
2. In the list of preset positions for the selected patrolling profile, select the preset position for which you want to change the time:

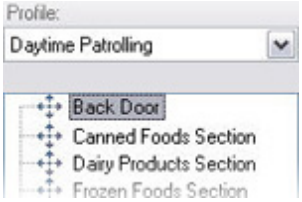

3. Specify the required time (in number of seconds) in the *Wait time (secs.)* field:

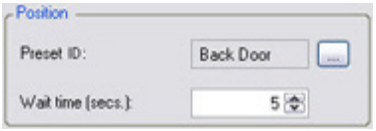

4. If required, repeat for other preset positions.

Customize transitions

By default, the time required for moving the camera from one preset position to another, known as *transition*, is estimated to be 3 seconds. During this time, motion detection is by default disabled on the camera, as irrelevant motion is otherwise likely to be detected while the camera moves between the preset positions. Transitions are also known as PTZ scanning.

Customizing speed for transitions is only supported if your camera supports PTZ scanning and is of the type where preset positions are configured and stored on your system's server (type 1 PTZ camera). Otherwise the *Speed* slider is grayed out.

You can customize the transitions between each of the preset positions in a patrolling profile. You are able to customize the following:

- The estimated transition time
- The speed with which the camera will move during a transition
- Which plug-ins to disable during transition.

To customize transitions between preset positions in a patrolling profile, do the following:

1. Select the required patrolling profile in the *Profile* list.
2. Select the *Customize transitions* check box:

This will add transition indications to the list of preset positions for the selected patrolling profile.

3. In the list, select the required transition:

4. Specify the estimated transition time (in number of seconds) in the *Expected time (secs.)* field:

5. Use the *Speed* slider to specify the required transition speed. When the slider is in its rightmost position, the camera will move with its default speed. The more you move the slider to the left, the slower the camera will move during the selected transition.

   Tip: A number indicating the exact speed is displayed near the top right corner of the slider. When required, the number (from 0.0001 (very slow) to 1.0000 (default speed)) allows you to define exactly the same custom speed across transitions.

6. In the *Plug-ins to disable* list, specify any plug-ins you want to disable during the selected transition. By default, the plug-in used for motion detection on the camera (*MotionDetectionPlugin*) is disabled in order to avoid irrelevant motion being detected during transition.

   To add a plug-in to the list, click *Add...*, and select the required plug-in. This requires that one or more other plug-ins are available, and that they can be disabled.

   To remove a plug-in from the list, for example if you do not want motion detection to be disabled during the transition, select the plug-in and click *remove*.

7. Repeat as required for other transitions.

Specify an end position

You are able to specify that the camera should move to a specific preset position when patrolling according to the selected patrolling profile ends.

1. Select the required patrolling profile in the *Profile* list.
2. Select the *Go to specific preset on finish* check box. This opens the *Select Preset* dialog.
3. In the *Select Preset* dialog, select the required end position, and click *OK*.

   Tip: You can select any of the camera's preset positions as the end position, you are not limited to the preset positions used in the patrolling profile.

4. The selected end position is added to the list of preset positions for the selected patrolling profile. When patrolling according to the selected patrolling profile ends, the camera will go to the specified end position.

Specify manual PTZ session timeout

Patrolling of PTZ cameras may be interrupted manually by Ocularis Client users with the necessary user rights.

You may specify how much time should pass before regular patrolling is resumed after a manual interruption:

1. In the Management Client's menu bar, select *Tools* > *Options*. This opens the *Options* window.
2. On the *Options* window's *General* tab, select the required amount of time in the *PTZ manual session timeout* list (default is 15 seconds).

The setting applies for all PTZ cameras on your system.

PTZ PRESETS TAB (CAMERA PROPERTIES)

The *Presets* tab lets you create preset positions to be used, for example, in rules (see "Manage rules" on page 165) for making a PTZ (Pan/Tilt/Zoom) camera move to a specific preset position when an event occurs, as well as in patrolling (see "PTZ Patrolling tab (camera properties)" on page 93), the automatic movement of a PTZ camera between a number of preset positions.

You manage preset positions on the *Presets* tab, which is available only when the selected camera is a PTZ (Pan/Tilt/Zoom) camera. The *Presets* tab will not be available if the selected PTZ camera does not support preset positions.

In order to use PTZ cameras' features, including the ability to pan, tilt, and zoom, operators must have a role which gives them the necessary rights. Refer to About roles (on page 182) for more information, including step-by-step descriptions of how to assign users to roles and how to specify the rights of roles.

Presets tab, with eight preset positions defined

Add a preset position (type 1)

As an alternative to defining preset positions in the system, preset positions may for some PTZ cameras also be defined on the camera device itself (typically by accessing a device-specific configuration web page) and imported into the system by selecting *Use presets from device* (see "*Use preset positions from device (type 2)*" on page 99).

To add a preset position for the camera in the system, do the following:

1. Click *Add...*. This will open the *Add Preset* window:

2. The *Add Preset* window displays a preview image from the camera; use the navigation buttons and/or sliders to move the camera to the required preset position. While you do this, you are able to verify the position of the camera through the preview image.
3. Specify a name or number for the preset position in the *Name* field.

   Tip: Use a descriptive name; the name may later be used in situations where you will not have access to details about the item, e.g. when using it in a rule.

4. Optionally, type a description of the preset position in the *Description* field.
5. Click *OK*. This will close the *Add Preset* window, and add the preset position to the *Presets* tab's list of available preset positions for the camera.

How to use the navigation buttons

The navigation buttons let you move the camera as follows:

Icon	What it does...	
	moves the view	up and to the left
		up
		up and to the right
		to the left
		to its default position
		to the right
		down and to the left
		down
		down and to the right
	Zooms in (one zoom level per click)	
	Zooms in (one zoom level per click)	

How to use the axes navigation sliders

The navigation sliders let you to move the camera along each of its axes. Click inside the sliders to move the sliders' red handles to the required positions.

The slider for the ***X-axis*** (allowing you to pan left/right) is located immediately below the preview image.

The slider for the ***Y-axis*** (allowing you to tilt the camera up/down) is located immediately to the left of the preview image.

The slider for the ***Z-axis*** (allowing you to zoom in and out) is located immediately above the preview image. The camera will zoom in when you move the slider towards *Tele*, and zoom out when you move the slider towards *Wide*.

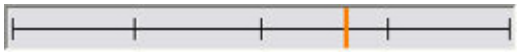

Example: *Add Preset* window's X-axis slider

How to use the iris slider

Iris settings are only available for some cameras.

Iris settings control the amount of light in images. The higher the iris setting, the lighter images will appear.

Click inside the slider to move the slider's red handle to the required position.

How to use the focus slider

Focus settings are only available for some cameras.

Click inside the slider to move the slider's red handle to the required position.

Use preset positions from device (type 2)

As an alternative to specifying preset positions in the system, preset positions may for some PTZ cameras also be defined on the camera device itself (typically by accessing a device-specific configuration web page.

Such device-defined presets can subsequently be imported into the system by selecting *Use presets from device*.

If importing presets from the camera device, any presets you have previously defined for the camera in will be removed; this will affect any patrolling profiles in which these presets are used, as well as any rules in which the affected patrolling profiles are used.

If you later wish to edit such device-defined presets, editing should take place on the camera device.

Assign a default preset position

If required, you are able to assign one of a PTZ camera's preset positions at the camera's default preset position.

Having a default preset position can be useful because it allows you to define rules (see "Manage rules" on page 165) specifying that the PTZ camera should go to the default preset position under particular circumstances, for example after the PTZ camera has been operated manually.

To assign a preset position as the default, select the required preset in your list of defined preset positions, then select the *default preset* box below the list.

Only one preset position can be the default preset position.

Edit a preset position

To edit an existing preset position defined in the system (presets imported from a device should be edited on the device itself), do the following:

1. Select the required preset position in the *Presets* tab's list of available preset positions for the camera.
2. Click *Edit...*. This opens the *Edit Preset* window:

Example only; features are camera-dependent

3. The *Edit Preset* window displays a preview image from the preset position in question; use the navigation buttons and/or sliders to change the preset position as required.
4. Change the name/number and description of the preset position as required.

   Tip: Use a descriptive name; the name may later be used in situations where you will not have access to details about the item, e.g. when using it in a rule.
5. Click *OK*.

Test a preset position

1. Select the required preset position in the *Presets* tab's list of available preset positions for the camera.
2. Click *Test*.
3. The *Presets* tab's preview image moves to the selected preset position.

   Tip: If the preview image does not appear to move to the selected preset position when you click *Test*, verify that preview image does not already show the selected preset position. In that case, try testing another preset position first.

360° Lens tab (Camera properties)

Use of 360° technology requires a dedicated ImmerVision 360° lens mounted.

In this way, 360° technology enables a.o.t. panomorph technology through an advanced lens.

The *360° Lens* tab lets you enable and configure panomorph support for the selected camera.

Tip: If you find that a camera with an ImmerVision 360° lens mounted takes a very long time to initialize, try turning the lens back and forth a bit. This long initializing time might be due to the lens no being fitted optimally.

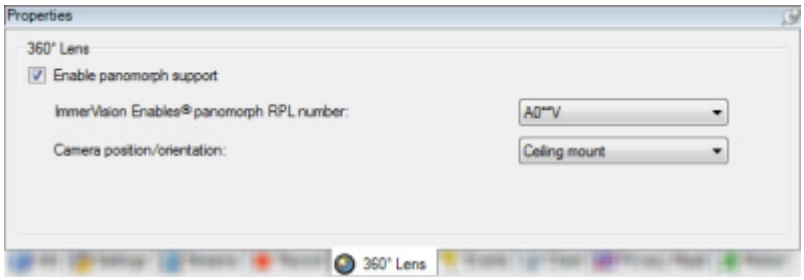

Enable and disable panomorph support

The panomorph feature is disabled by default. To enable/disable it, select/clear the *360° Lens* tab's *Enable panomorph support* check box.

Panomorph settings

When enabling the panomorph support functionality, you must also select a Registered Panomorph Lens (RPL) number from the *ImmerVision® Enables panomorph RPL number* list. This is to ensure identification and correct configuration of the lens used with the camera in question. The RPL number is usually found on the lens itself or on the box it came in. For details of ImmerVison, panomorph lenses, and RPLs, see http://www.immervision.com/en/home/index.php**.**

You must also indicate the physical position/orientation of the camera in question. This is done by selecting its position from the *Camera position/orientation* list.

STREAMS TAB (CAMERA PROPERTIES)

Depending on the recording component, functionality described here may be limited or unavailable.

To access the *Streams* tab, expand *Devices* in the Management Client's Site navigation pane (see "Panes Overview" on page 33), expand the relevant camera folder in the Overview pane (see "Panes Overview" on page 33), select the required camera and then select the *Streams* tab in the Properties pane (see "Panes Overview" on page 33).

The tab will by default list a single stream—the selected camera's default stream, used for live video as well as for video which is being recorded for playback purposes.

Note that while it is possible to set up and use two live streams, only one of the enabled live streams is able to record video at a time. To change which stream to use for recording, use the *Record* box.

About multi-streaming (on page 101)

Add a new stream (see "Add a stream" on page 102)

About multi-streaming

You manage multi-streaming on the *Streams* tab. The tab is only available when the selected camera or device group supports multi-streaming.

Viewing of live video and playing back of recorded video does not necessarily require the same settings to achieve the best result. To handle this, your system and some cameras support multi-streaming, with which you can establish two independent streams to the recording server. **Either** one stream for live viewing and another stream for playback purposes **or** two separate live streams—with different resolution, encoding, and frame rate.

Example 1, live and recorded video:

- For viewing **live** video, your organization may prefer MPEG4 at a high frame rate.

- For playing back **recorded** video, your organization may prefer MJPEG at a lower frame rate because this will help preserve disk space.

Example 2, two live videos:

- For viewing **live video from a local operating point**, your organization may prefer MPEG4 at a high frame rate to have the highest quality of video available.
- For viewing **live video from a remotely connected operating point**, your organization may prefer MJPEG at a lower frame rate and quality in order to preserve network bandwidth.

Even when cameras support multi-streaming, individual multi-streaming capabilities may vary considerably between different cameras. Refer to camera's documentation for exact information. To see if a camera offers different types of streams, refer to the *Settings* tab (see "Settings tab overview" on page 121). Also note that cameras under remote sites in an OnSSI Interconnect setup (see "About OnSSI Interconnect" on page 57) only support single streams.

If you select a device group with 400 or more cameras, the *Streams* tab will not be available for viewing and editing because changing settings for so many devices in one go takes too long time.

Add a stream

1. On the *Streams* tab, click *Add*. This will add a second stream to the list (you cannot have more than two streams).
2. (Optional) In the *Name* column, edit the name of the stream.
3. In the *Live Mode* column, select when live streaming is needed.
4. In the *Default* column, select which stream is the default one.
5. In the *Record* column, select the check box if you want to use the stream for recorded video or leave it cleared if you only want to use it for live video.
6. In the *Edge Recording* column it is indicated whether the selected stream supports edge recording (see "Record tab overview" on page 125) or not.
7. Click *Save.*

Manage Microphones

On many devices you are able to attach external microphones and some devices even have built-in microphones.

Devices' microphones are automatically detected when you add the devices to your system through the Management Client's *Add Hardware* (on page 53) wizard, regardless of which of the wizard's detection options you use.

Microphones do not require separate licenses; you can use as many microphones as required on your system.

You can use microphones completely independently of cameras.

Who is able to listen to audio recorded by microphones? Users of the Ocularis Client can—provided microphones are available, and the users have the rights to use them—listen to audio from microphones. Roles determine users' right to listen to microphones. You cannot listen to microphones from the Management Client.

Tip: the system comes with a default rule which ensures that audio feeds from all connected microphones and speakers are automatically fed to the system. Like other rules, the default rule can be deactivated and/or modified as required.

You have two entry points for managing microphones:

- In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Devices*, select *Microphones*, expand the required device group, and select the required microphone. If no device groups are available, you must first group your microphones. Refer to About Device Groups (on page 104) for information about creating groups as well as adding microphones to your groups.

- In the Management Client's Site Navigation pane, expand *Servers* and select *Recording Servers*, then in the Overview pane (see "Panes Overview" on page 33) expand the required recording server, expand the required device and select the required microphone.

Check the product release notes to verify that microphones are supported for the devices and firmware used.

ENABLE MICROPHONES

When microphones are detected with the wizard *Add Hardware* (on page 53) they are by default disabled. You can enable microphones when needed. If a device has several microphones you can enable one, some, or all of them as required.

1. In the Site Navigation pane (see "Panes Overview" on page 33), expand *Servers* and select *Recording Servers*.
2. In the Overview pane (see "Panes Overview" on page 33), expand the relevant recording server, and find the device on which the microphone is placed.
3. Right-click the required microphone, and select *Enabled*.

On some devices, a microphone can also be enabled/disabled on the device itself, typically through the device's own configuration web page. If a microphone does not work after enabling it in the Management Client, you should verify whether the problem may be due to the microphone being disabled on the device itself.

CONFIGURE INDIVIDUAL MICROPHONES

You configure individual microphones by selecting the required microphone in the list, then specifying the microphone's required settings on the tabs in the Properties pane (see "Panes Overview" on page 33):

Tab	Use for specifying
Info (see "***Info tab overview***" on page 119)	The selected microphone's name, etc.
Settings (see "***Settings tab overview***" on page 121)	The selected microphone's general settings.
Record (see "***Record tab overview***" on page 125)	The selected microphone's recording, database and archiving storage settings.
Events (see "***Events tab overview***" on page 128)	Events.

VIEW CURRENT STATE OF MICROPHONES

When you have selected a microphone in the Management Client, information about the current status of the selected microphone is presented in the Preview pane (see "Panes Overview" on page 33).

When the microphone is not active, it is shown as:

When the microphone is active, it is shown as:

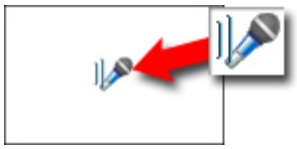

READ MICROPHONE LIST'S STATUS ICONS

The following icons are used to indicate status of cameras (see "Manage cameras" on page 82), microphones (see "Manage Microphones" on page 102), speakers (see "Manage speakers" on page 107), input (see "Manage input" on page 109) and output (see "Manage output" on page 113) events in item lists:

Cam-era	Micro phone	Spea-ker	In-put	Out-put	Description
					**Item enabled:** Can communicate with the recording server, and can if required be started/stopped automatically through a rule.
					**Item recording.** **Speaker being recorded:** Note that what is being said through the speaker can be recorded, but cannot subsequently be played back or exported (for example to prove that a warning was given).
					**Item temporarily stopped or has no feed:** Often shown when an item is communicating with the system while it is being disabled or enabled. Also shown if the Default Start Audio Feed Rule is not active; see Managing Rules. When stopped, no information is transferred to the system. In which case—if it is a camera—neither live viewing nor recording will be possible. However, a stopped item will still be able to communicate with the recording server, as opposed to when an item is disabled.
					**Item disabled:** Cannot be started automatically through a rule and will not be able to communicate with the recording server. In the case of a camera, when a camera is disabled, neither live viewing nor recording will be possible.
					**Item database being repaired.**
					**Item requires attention.**
					**Status unknown.**
					Note that some icons may be combined, as in this example where **Item is enabled** is combined with **Item is recording** (since a recording item is also an enabled item).

About device groups

You are able to group different types of devices (cameras, microphones, speakers, inputs, outputs) on your system by using device groups. The use of device groups has several benefits:

- Device groups help you maintain an intuitive overview of devices on your system

- You are able to specify common properties for all devices within a device group in one go
- When dealing with roles (see "About roles" on page 182), you are able to specify common security settings for all devices within a device group in one go
- When dealing with rules (see "Manage rules" on page 165), you are able to apply a rule for all devices within a device group in one go

You can add as many device groups as required; you are completely free to decide which devices to include. The only restriction is that you cannot mix different types of devices (for example cameras and speakers) in a device group.

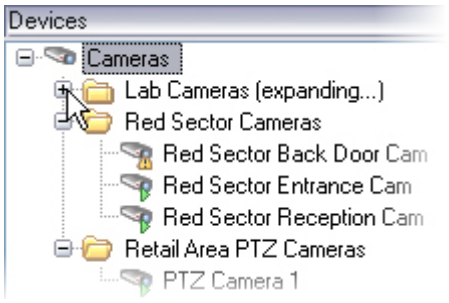

Example: cameras grouped into device groups

If a device group contains 400 devices or more, the *Settings* tab is unavailable for viewing and editing. For camera groups, the *Streams* tab is also unavailable for editing and viewing if the group contains 400 cameras or more. When you click the plus sign next to the device folder, your system will load the contents of the device folder, which may take a few seconds. While expanding, the text *(expanding...)* is displayed next to the folder name.

Note that if you delete a device group, you only delete the device group itself. If you wish to delete IP hardware (see "About hardware" on page 55) - such as a camera - from your system, do so on a recording server level.

The following examples are based on grouping cameras into device groups, but the principle applies for microphones, speakers, inputs and outputs as well.

ADD A DEVICE GROUP

1. In the Overview pane (see "Panes Overview" on page 33), right-click the item under which you wish to create the new device group.
2. Select *Add Device Group*:

The *Add Device Group* dialog will appear.

3. In the *Add Device Group* dialog, specify a name and description of the new device group:

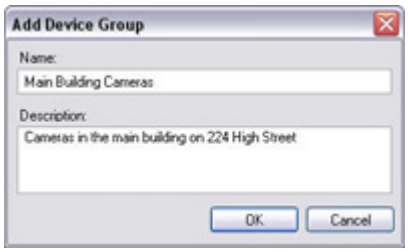

The description will later appear when pausing the mouse pointer over the device group in the device list.

4. Click *OK*. A folder representing the new device group is added to the list. You are now able to specify which devices should belong in the device group.

   Tip: If required, you are able to add device groups as subgroups under other device groups, as illustrated here:

SPECIFY WHICH DEVICES TO INCLUDE IN A DEVICE GROUP

1. In the Overview pane (see "Panes Overview" on page 33)'s device list, right-click the required device group folder.
2. Select *Edit Device Group Members.* The *Select Group Members* window will appear.
3. In the *Select Group Members* window, select either:
   - The *Device Groups* tab, which lists devices based on existing device groups.

     Tip: A device can be a member of more than one device group.

   - The *Recording Servers* tab, which lists devices based on which recording servers the devices belong on.
4. Select the devices you wish to include, and click *Add*. This will copy the selected devices to the *Selected* box:

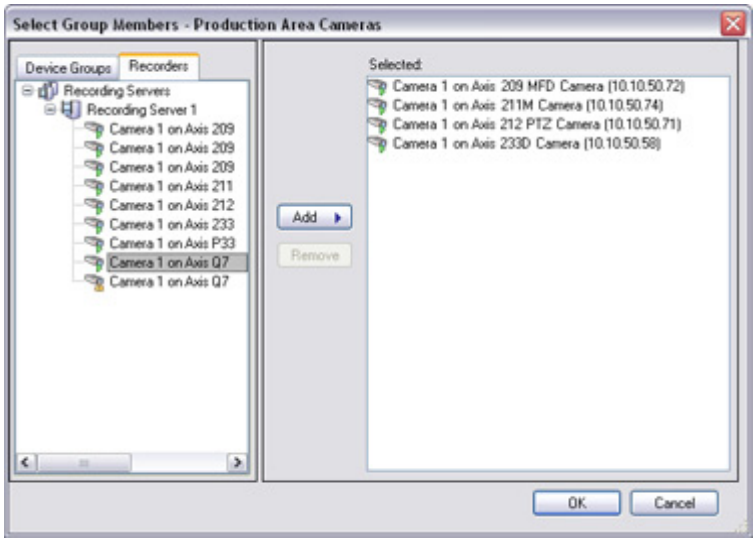

Tip: You may also double-click a device to copy it from one box to the other, or you may drag devices between the two boxes.

Tip: To select several devices in one go, press the CTRL key on your keyboard while selecting.

5. Click *OK*. The selected devices will be added to your device group on the device list.

SPECIFY COMMON SETTINGS FOR ALL DEVICES IN A DEVICE GROUP

When using device groups, you are able to quickly specify common properties for all devices within a given device group:

1. In the Overview pane (see "Panes Overview" on page 33)'s device list, click the required device group.

   In the Properties pane (see "Panes Overview" on page 33), all properties *which are available on all of the device group's devices* will be listed, grouped on tabs.

2. Specify the required common properties.

   Properties not available on all *of the devices in the device group will not be listed; such properties must still be configured individually for each device.*

   If the device group contains 400 or more devices the Settings *tab (see "Remote recording - camera/remote system" on page 128) is unavailable for viewing and editing. For camera groups the* Streams *tab is also unavailable for viewing and editing if the group contains 400 cameras or more.*

   Tip: The Settings tab has convenient functionality for quickly switching between settings for the device group and settings for individual devices.

Manage speakers

On many devices you are able to attach external loudspeakers; some devices even have built-in speakers.

Devices' speakers are automatically detected when you add the devices to your system through the Management Client's *Add Hardware* (on page 53) wizard, regardless of which of the wizard's detection options you use. Speakers do not require separate licenses; you can use as many speakers as required on your system.

You can use speakers completely independently of cameras.

Who is able to talk through speakers? Users of the Ocularis Client can—provided speakers are available, and the users have the rights to use them—click a button to talk through speakers. Roles determine users' right to talk through speakers. You cannot talk through speakers from the Management Client.

What happens if two users want to speak at the same time? Roles determine users' right to talk through speakers. As part of the roles definition, you are able to specify a speaker priority from very high to very low. If two users want to speak at the same time, the user whose role has the highest priority will win the ability to speak. If two users with the same role want to speak at the same time, the first-come first-served principle applies.

Tip: Your system comes with a default rule which ensures that audio feeds from all connected microphones and speakers are automatically fed to the system. Like other rules, the default rule can be deactivated and/or modified as required.

You have two entry points for managing speakers:

- In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Devices*, select *Speakers*, expand the required device group, and select the required speaker. If no device groups are available, you must first group your speakers. Refer to About Device Groups (on page 104) for information about creating groups as well as adding speakers to your groups.
- In the Management Client's Site Navigation pane, expand *Servers* and select *Recording Servers.* In the Overview pane (see "Panes Overview" on page 33), expand the required recording server and select the required speaker.

Check the product release notes to verify that speakers are supported for the devices and firmware used.

ENABLE SPEAKERS

When speakers are detected with the wizard *Add Hardware*, they are by default disabled. You can enable speakers when needed. If a device has several speakers you can enable one, some, or all of them as required.

1. In the Site Navigation pane (see "Panes Overview" on page 33), expand *Servers* and select *Recording Servers*.
2. In the Overview pane (see "Panes Overview" on page 33), expand the relevant recording server, and find the device on which the speaker is placed.
3. Right-click the required speaker, and select *Enabled*.

On some devices, a speaker can also be enabled/disabled on the device itself, typically through the device's own configuration web page. If a speaker does not work after enabling it in the Management Client, you should verify whether the problem may be due to the speaker being disabled on the device itself.

CONFIGURE SPEAKERS

You configure individual speakers by selecting the required speaker in the list, then specifying the speaker's required settings on the tabs in the Properties pane (see "Panes Overview" on page 33):

Tab	Use for specifying
Info (see "Info tab overview" on page 119)	The selected speaker's name, etc.
Settings (see "Settings tab overview" on page 121)	The selected speaker's general settings.
Record (see "Record tab overview" on page 125)	The selected speaker's recording, database and archiving storage settings.

VIEW CURRENT STATE OF SPEAKERS

When you have selected a speaker in the Management Client, information about the current status of the selected speaker is presented in the Preview pane (see "Panes Overview" on page 33).

When a speaker is not active, it is shown as:

When a speaker is active, it is shown as:

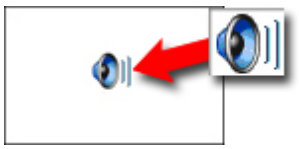

READ SPEAKER LIST'S STATUS ICONS

The following icons are used to indicate status of cameras (see "Manage cameras" on page 82), microphones (see "Manage Microphones" on page 102), speakers (see "Manage speakers" on page 107), input (see "Manage input" on page 109) and output (see "Manage output" on page 113) events in item lists:

Cam-era	Micro phone	Spea-ker	In-put	Out-put	Description
					**Item enabled:** Can communicate with the recording server, and can if required be started/stopped automatically through a rule.
					**Item recording.** **Speaker being recorded:** Note that what is being said through the speaker can be recorded, but cannot subsequently be played back or exported (for example to prove that a warning was given).
					**Item temporarily stopped or has no feed:** Often shown when an item is communicating with the system while it is being disabled or enabled. Also shown if the Default Start Audio Feed Rule is not active; see Managing Rules. When stopped, no information is transferred to the system. In which case—if it is a camera—neither live viewing nor recording will be possible. However, a stopped item will still be able to communicate with the recording server, as opposed to when an item is disabled.
					**Item disabled:** Cannot be started automatically through a rule and will not be able to communicate with the recording server. In the case of a camera, when a camera is disabled, neither live viewing nor recording will be possible.
					**Item database being repaired.**
					**Item requires attention.**
					**Status unknown.**
					Note that some icons may be combined, as in this example where **Item is enabled** is combined with **Item is recording** (since a recording item is also an enabled item).

Manage input

On many devices you are able to attach external units to input ports on the device. Input units are typically external sensors. Such external sensors may, for example, be used for detecting if doors, windows, or gates are opened. Input from such external input units is treated as events by the system.

Such events can be used in rules (see "Manage rules" on page 165). For example, you could create a rule specifying that a camera should begin recording when an input is activated, and stop recording 30 seconds after the input is deactivated.

Devices' input ports are automatically detected when you add the devices to your system through the Management Client's Add Hardware (on page 53) wizard, regardless of which of the wizard's detection options you use.

You have two entry points for managing input:

- In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Devices*, select *Inputs*, expand the required device group, and select the required input. If no device groups are available, you must first group your input. Refer to About Device Groups (on page 104) for information about creating groups as well as adding input to your groups.

- In the Management Client's Site Navigation pane, expand *Servers* and select the *Recording Server* node, then expand the required recording server in the Overview pane (see "Panes Overview" on page 33) and select the required input.

Before you specify use of external input and output units on a device, verify that sensor operation is recognized by the device. Most devices are capable of showing this in their configuration interfaces, or via Common Gateway Interface (CGI) script commands. Also check the system release notes to verify that input- and output-controlled operations are supported for the devices and firmware used.

ENABLE INPUT

When inputs are detected with the *Add Hardware* (on page 53) process, they are by default disabled. You can activate inputs when needed. If a device has several inputs you can enable one, some, or all of them as required.

1. In the Site Navigation pane (see "Panes Overview" on page 33), expand *Servers* and select *Recording Servers.*
2. In the Overview pane (see "Panes Overview" on page 33) expand the relevant recording server, and find the device on which the input is placed.
3. Right-click the required input, and select *Enabled*.

SPECIFY INPUT PROPERTIES

Each input typically has several properties. You can access these properties in two ways:

- In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Devices,* and select *Inputs*. In the Overview pane (see "Panes Overview" on page 33) expand the required inputs folder and select the required input.

  - or -

- In the Overview pane, select a device group to define settings for all inputs in the group, or expand a device group, and select the required input.

The properties of the selected input, or the common properties for all inputs in a selected device group, will be displayed on the following tabs: *Settings*, *Info*, and *Events*.

VIEW THE CURRENT STATE OF AN INPUT

The change of an input's state is regarded as an event by the system. Events can be used in rules and hereby trigger actions when the state of an input is changed.

Refer to Define input- and output-related rules (see "Define in- and output-related rules" on page 117) for more information about how to include an input event in a rule.

To view the current state of an input in the Management Client, do the following:

1. In the Site Navigation pane (see "Panes Overview" on page 33), expand *Devices*, and select *Inputs*.
2. In the Overview pane (see "Panes Overview" on page 33), expand the required inputs folder and select the required input.

   Tip: You may select a group of inputs to view the current status of all inputs in the group.

3. Information about the current status of the selected input is presented in the Preview pane (see "Panes Overview" on page 33).

   When an input is deactivated, it is shown by a gray indicator:

When the input is activated, the indicator lights up green:

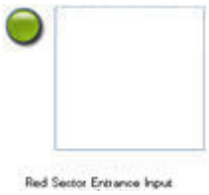

FILL IN PROPERTIES ON THE INFO TAB

Lets you view and edit basic information about an input. Contains the following fields:

Name	Description
Name	Name of the input. Optional, but highly recommended. Used whenever the input is listed in the system and clients. Does not have to be unique. To change the name, overwrite the existing name and click *Save* in the toolbar (see "Management Client Overview" on page 30). **Tip:** If you change the name, it will be updated throughout the system. This means that if the name is used in, for example, a rule, the name will automatically change in the rule as well.
Description	Description of the input. Optional. Will appear in a number of listings within the system. For example, the description will appear when pausing the mouse pointer over the item's name in the Overview pane (see "Panes Overview" on page 33): 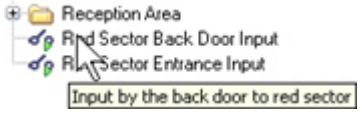 To specify a description, type the description and click ***Save*** in the toolbar (see "Management Client Overview" on page 30).
Hardware name	Name of the hardware with which the input unit is connected. The field is non-editable from here, but can be changed by clicking *Go To* next to it. This takes you to hardware information, where the name is editable.
Unit number	Non-editable field, displaying the unit on which the input can be found on the hardware. For hardware capable of having more than one input unit attached, the unit number will typically indicate the number of the input port to which the input is attached. For hardware with, for example, four input ports, the numbers will typically range from *0* to *3*.

FILL IN SETTINGS TAB PROPERTIES

The content of the *Settings* tab is determined entirely by the devices in question, and may vary depending on the input selected. Verify or edit key input settings, for a selected input, or for all inputs within a selected device group. If the selected device group contains 400 or more inputs, the *Settings* tab will be unavailable for viewing and editing because changing settings for so many devices in one go takes too long time.

Content may vary, but you typically see the following property:

Name	Description

Input rises on	Define whether the input signal should be considered rising on *Circuit closed* or *Circuit open*. The value of this setting is used on the input's *Events* tab, where you define properties for input events: *Input Rising* event, *Input Falling* event, and *Input Changed* event. See also the description of the properties of the ***Events*** tab (see "Fill in properties on the Events tab" on page 112).

The content of the *Settings* tab is displayed in a table, in which the first column lists the available settings, and the second column lists the value of each setting. You can change the values of these settings if needed.

FILL IN PROPERTIES ON THE EVENTS TAB

Lets you define events based on changes of the input's state, from circuit opened to circuit closed or the reverse order. The events you define can subsequently be used in rules.

You can define events for a selected input, but not for all inputs in a device group.

1. In the Site Navigation pane (see "Panes Overview" on page 33), expand *Devices*, and select *Inputs*.
2. In the Overview pane (see "Panes Overview" on page 33), select the required input.
3. Select the *Events* tab, and click *Add...*.

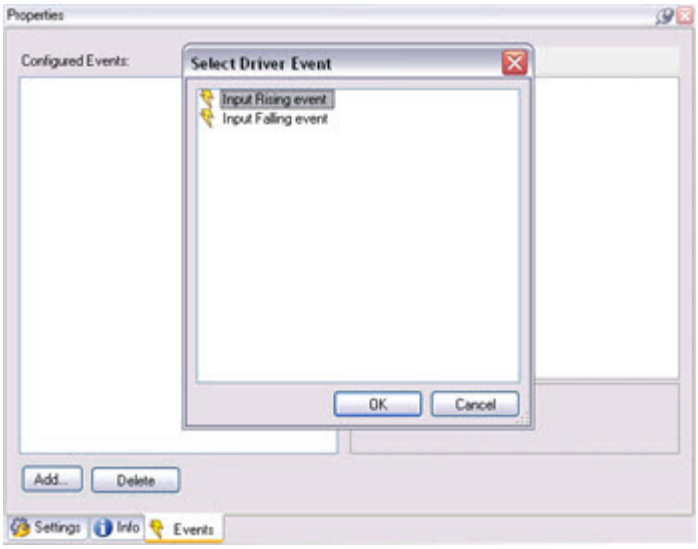

4. In the *Select Driver Event* dialog, select the appropriate option (*Input Rising event*, *Input Falling event*, or *Input Changed event*).
5. Click *OK*. Your selected type of input event will now appear in the *Events* tab's *Configured events* list.

To the right of the list, settings for the selected input event are displayed in a table. The table's first column lists available settings, the second column lists the value of each setting.

The settings on the *Events* tab is determined entirely by the relevant devices, and is likely to vary depending on the input selected.

Content may vary, but you will typically see the following property:

- ***Enabled****:* Select between *True* (enabled), or *False* (disabled).

You are typically able to change the values:

1. Select the row with the property you want to change.
2. Click the  button to the right of the properties column.

3. Change the value of the property.
4. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

When you have changed a setting to a non-default value, the value will appear in **bold**. When a value must be within a certain range, for example between 0 and 100, the allowed range will be displayed in the gray information box below the settings table.

READ THE INPUT LIST'S STATUS ICONS

The following icons are used to indicate status of cameras (see "Manage cameras" on page 82), microphones (see "Manage Microphones" on page 102), speakers (see "Manage speakers" on page 107), input (see "Manage input" on page 109) and output (see "Manage output" on page 113) events in item lists:

Cam-era	Micro phone	Spea-ker	In-put	Out-put	Description
					**Item enabled:** Can communicate with the recording server, and can if required be started/stopped automatically through a rule.
					**Item recording.** **Speaker being recorded:** Note that what is being said through the speaker can be recorded, but cannot subsequently be played back or exported (for example to prove that a warning was given).
					**Item temporarily stopped or has no feed:** Often shown when an item is communicating with the system while it is being disabled or enabled. Also shown if the Default Start Audio Feed Rule is not active; see Managing Rules. When stopped, no information is transferred to the system. In which case—if it is a camera—neither live viewing nor recording will be possible. However, a stopped item will still be able to communicate with the recording server, as opposed to when an item is disabled.
					**Item disabled:** Cannot be started automatically through a rule and will not be able to communicate with the recording server. In the case of a camera, when a camera is disabled, neither live viewing nor recording will be possible.
					**Item database being repaired.**
					**Item requires attention.**
					**Status unknown.**
					Note that some icons may be combined, as in this example where **Item is enabled** is combined with **Item is recording** (since a recording item is also an enabled item).

Manage output

On many devices you are able to attach external units to output ports on the device. This allows you to activate/deactivate lights, sirens, etc. through the system.

Output may be used when creating rules (see "Manage rules" on page 165). You can create rules that automatically activate or deactivate outputs, and rules that trigger actions when the state of an output is changed.

Output can also be triggered manually from the Management Client and the **Ocularis Client**.

Devices' output ports are automatically detected when you add the devices to the system through the Management Client's *Add Hardware* (on page 53) wizard, regardless of which of the wizard's detection options you use. By default, output are disabled. You can enable output when needed.

You have two entry points for managing outputs:

- In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand **Devices**, select **Outputs**, expand the required device group, and select the required output. If no device groups are available, you must first group your output. Refer to About Device Groups (on page 104) for information about creating groups as well as adding output to your groups.
- In the Management Client's Site Navigation pane, expand **Servers** and select **Recording Servers**, then in the Overview pane (see "Panes Overview" on page 33) expand the required recording server and select the required output.

Before you specify use of external input and output units on a device, verify that sensor operation is recognized by the device. Most devices are capable of showing this in their configuration interfaces, or via Common Gateway Interface (CGI) script commands. Also check the system release notes to verify that input- and output-controlled operations are supported for the devices and firmware used.

ENABLE OUTPUT

When outputs are detected with the *Add Hardware* (on page 53) process, they are by default disabled. You can activate outputs when needed.

If a device has several outputs you can enable one, some, or all of them as required.

1. In the Site Navigation pane (see "Panes Overview" on page 33), expand *Servers* and select *Recording Servers*.
2. In the Overview pane (see "Panes Overview" on page 33) select the relevant recording server, and find the device on which the output is placed.
3. Right-click the required output, and select *Enabled*.

SPECIFY OUTPUT PROPERTIES

Each output has several properties which can be found on the output's *Settings* and *Info* tabs. You can access these tabs in two ways:

- In the Site Navigation pane (see "Panes Overview" on page 33), expand *Devices* and select *Outputs*, then in the Overview pane (see "Panes Overview" on page 33) expand the required outputs folder and select the required output.

  - or -

- In the Overview pane, select a device group to change the settings for all outputs in this group, or expand a device group and select the required output.

The properties of the selected output, or the common properties for all outputs in a selected device group, will be displayed on the following tabs: *Settings* and *Info*.

AUTOMATIC/MANUAL ACTIVATION OF OUTPUT

Output can be activated automatically or manually:

- **Automatic activation of output**

  With the Management Client's rules (see "Manage rules" on page 165) feature, you are able to create rules that automatically activate or deactivate output, and rules that trigger actions when the state of an output is changed.

For example, you may create a rule specifying that a siren should sound if motion is detected on a particular camera, or you may create a rule specifying that a camera should start recording if a siren sounds. Refer to Define Input- and Output-Related Rules (see "Define in- and output-related rules" on page 117) for more information.

- **Manual activation of output**

  Output may be activated manually from the Management Client and the Ocularis Client:

  1. In the Site Navigation pane (see "Panes Overview" on page 33), expand *Devices* and select *Outputs*.
  2. In the Overview pane (see "Panes Overview" on page 33), expand the required outputs folder and select the required output.

     Tip: You may select a group of outputs, for example *All Outputs*, to manually activate all outputs in the group.
  3. The availability of features for manually activating an output depends entirely on the device in question, and may vary.
  4. Typically, the following elements are shown for each output in the Preview pane (see "Panes Overview" on page 33):

     

  5. Select/clear the check box to activate/deactivate the selected output. When an output is activated, the indicator lights up green:

     

  6. Alternatively, click the rectangular button to activate the output for the duration defined in the *Output Trigger Time* setting on the *Settings* tab (this feature/setting may not be available for all outputs). After the defined duration, the output is automatically deactivated.

FILL IN PROPERTIES ON THE INFO TAB

Lets you view and edit basic information about an output:

Name	Description
Name	Name of the output. Optional, but highly recommended. Used whenever the output is listed in the system and clients. Does not have to be unique. To change the name, overwrite the existing name and click *Save* in the toolbar (see "Management Client Overview" on page 30). **Tip:** If you change the name, it will be updated throughout the system. This means that if the name is used in, for example, a rule, the name will automatically change in the rule as well.

Name	Description
Description	Description of the output. Optional. Will appear in a number of listings within the system. For example, the description will appear when pausing the mouse pointer over the item's name in the Overview pane (see "Panes Overview" on page 33): 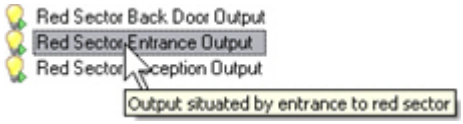 To specify a description, type the description and click ***Save*** in the toolbar (see "Management Client Overview" on page 30).
Hardware name	Name of the hardware with which the output unit is connected. The field is non-editable from here, but can be changed by clicking *Go To* next to it. This takes you to hardware information, where the name is editable.
Unit number	Non-editable field, displaying the unit on which the output can be found on the hardware. For hardware capable of having more than one output unit attached, the unit number will typically indicate the number of the output port to which the output is attached. For hardware with, for example, four output ports, the numbers will typically range from 0 to 3.

FILL IN PROPERTIES ON THE SETTINGS TAB

Lets you verify or edit key output settings, such as active output state, output trigger time, etc., for a selected output, or for all outputs within a selected device group. However, if the device group contains 400 cameras or more the *Settings* tab will not be available for viewing and editing because changing settings for so many devices in one go takes too long time.

The content of the *Settings* tab is determined entirely by the drivers for the cameras in question, and is likely to vary depending on the output selected.

Some devices are only able to apply outputs for a relatively short time, for example max. 5 seconds. Refer to the documentation for the device in question for exact information.

Content is displayed in a table, in which the first column lists the available settings, and the second column lists the value of each setting.

You are typically able to change the values:

1. Select the row with the property you want to change
2. Click the  button to the right of the properties column.
3. Change the value of the property.
4. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

When you have changed a setting to a non-default value, the value will appear in **bold**. When a value must be within a certain range, for example between 0 and 100, the allowed range will be displayed in the gray information box below the settings table.

READ THE OUTPUT LIST'S STATUS ICONS

The following icons are used to indicate status of cameras (see "Manage cameras" on page 82), microphones (see "Manage Microphones" on page 102), speakers (see "Manage speakers" on page 107), input (see "Manage input" on page 109) and output (see "Manage output" on page 113) events in item lists:

Cam-era	Micro phone	Spea-ker	In-put	Out-put	Description
					**Item enabled:** Can communicate with the recording server, and can if required be started/stopped automatically through a rule.
					**Item recording.** **Speaker being recorded:** Note that what is being said through the speaker can be recorded, but cannot subsequently be played back or exported (for example to prove that a warning was given).
					**Item temporarily stopped or has no feed:** Often shown when an item is communicating with the system while it is being disabled or enabled. Also shown if the Default Start Audio Feed Rule is not active; see Managing Rules. When stopped, no information is transferred to the system. In which case—if it is a camera—neither live viewing nor recording will be possible. However, a stopped item will still be able to communicate with the recording server, as opposed to when an item is disabled.
					**Item disabled:** Cannot be started automatically through a rule and will not be able to communicate with the recording server. In the case of a camera, when a camera is disabled, neither live viewing nor recording will be possible.
					**Item database being repaired.**
					**Item requires attention.**
					**Status unknown.**
					Note that some icons may be combined, as in this example where **Item is enabled** is combined with **Item is recording** (since a recording item is also an enabled item).

Define in- and output-related rules

To be able to automatically

- activate an output or trigger an event activated by an output, you must, after you have enabled an output,
- trigger an action activated by an input, you must, after you have enabled the input and created an event based on the input,

include it in a rule.

refer to Manage Output (on page 113) or Manage Input (on page 109) for more information.

For example, you may create a rule specifying that:

- a siren should sound if motion is detected on a particular camera, or you may create a rule specifying that a camera should start recording if a siren sounds (output).
- a camera should record if a particular input is activated (input).

For the following examples to be useful you should have general knowledge about managing rules (see "Manage rules" on page 165). Also refer to Create Typical Rules (on page 142) for other step-by-step descriptions of how to create rules.

Tip: When you create a rule based on an in- or output event, the actions you specify in the rule do not have to relate to the device on which the external in- or output was activated; you can easily specify that the actions should take place on one or more different devices— even across recording servers.

RULE THAT ACTIVATES/DEACTIVATES AN OUTPUT

1. Start the *Manage Rule* and in step 1 select a rule type and, if necessary, a condition in step 2.
2. In *Manage Rule*'s step 3 (*Step 3: Actions*) select the *Set device output to <state>* action.
3. If you like the output to be activated/deactivated immediately, skip this step. If you do not want to activate or deactivate the output immediately after the event, click the *immediately* link in the initial rule description, and select an interval between the event and the activation/deactivation of the output. Click *OK* to confirm your selection.
4. Click the *state* link in the initial rule description, and select whether you want to activate or deactivate the output. Click *OK* to confirm your selection.
5. Click the *devices* link in the initial rule description, and select which output you want to activate or deactivate. Click *OK* to confirm your selection.
6. If wanted you can select more actions in the *Manage Rul*e's step 3 (*Step 3: Actions*). Do so or click *Next* to continue to the next step.
7. In *Manage Rule*'s step 4 (*Step 4: Stop criteria*) select one of the stop actions, for instance to deactivate the output after a certain time or event.
8. Click *Finish* to save the rule.

RULE THAT MAKES AN OUTPUT TRIGGERS AN ACTION

In the *Rules* feature, all registered external output (activation, deactivation or change) is treated as an event. Based on an event, you are able to specify a wide variety of actions to take.

To define a rule where an output activates an action, do the following:

1. Start the *Manage Rule*. In step 1 (*Type of rule*) select *Perform an action on <event >*.
2. Click *event* in the initial rule description.
3. In the *Select an Event* dialog, in the ***Devices***, ***Predefined Events*** group (see "Predefined events, devices" on page 162), select the appropriate option for your rule: *Output Activated*, *Output Changed* or *Output Deactivated*. Click *OK*.
4. Click *devices/recorders/servers* in the initial rule description.
5. In the *Select Devices and Groups* dialog select the required output. Click *OK*.
6. Click *Next* to continue to step 2 (*Conditions*) and if needed select a condition. Continue to step 3 (*Actions*) and select one or more actions.
7. If you do not want to define a stop action, skip this step. If you want to define a stop action— for instance to deactivate the output again— click *Next* to continue to step 4 (*Stop criteria*), and select a stop action. Click *Finish*.

RULE THAT MAKES AN INPUT TRIGGER AN ACTION

In the *Rules* feature, all registered external input (activation, deactivation, or change) is treated as an event. Based on an event, you are able to specify a wide variety of actions to take.

To define a rule specifying that an input should result in one or more actions (for example the starting of recording on a certain camera), do the following:

1. Start *Managing Rules*. In step 1 (*Type of rule*), select *Perform an action on <event >*.
2. Click *event* in the initial rule description.
3. In the *Select an Event* dialog, in the ***Devices***, ***Configurable Events*** group (see "Configurable events, devices" on page 162), select the appropriate option for your rule: *Input Activated*, *Input Changed*, or *Input Deactivated*. Click *OK*.
4. Click *devices/recording servers/management servers* in the initial rule description.
5. In the *Select Devices and Groups* dialog select the required input. Click *OK*.
6. Continue to step 2 (*Conditions*) and, if needed, select a condition. Continue to step 3 (*Actions*) and select one or more actions. Continue to step 4 (*Stop criteria*), and select a stop criteria. Continue to step 5 (*Stop actions*), and select a stop action. Click *Finish*.

Info tab overview

The *Info* tab lets you view and edit basic information about a selected item in a number of fields. The following items under *Devices* have an *Info* tab:

- Cameras (see "Manage cameras" on page 82)
- Hardware (see "About hardware" on page 55)
- Microphones (see "Manage Microphones" on page 102)
- Speakers (see "Manage speakers" on page 107)

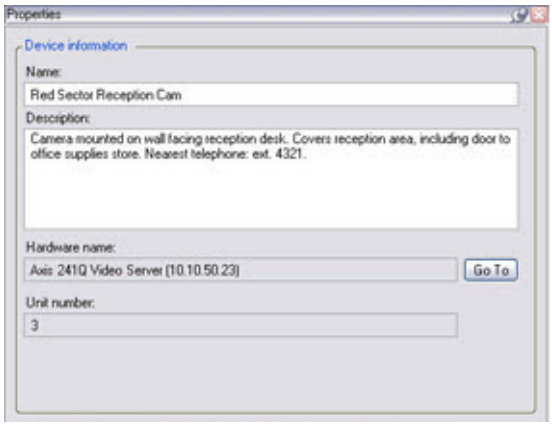

Example of *Info* tab from a camera...

SPECIFY HARDWARE AND DEVICE INFO PROPERTIES

Name	Description
Name	Name of the item. Not compulsory, but highly recommended. Used whenever the item is listed in the system and clients. Does not have to be unique. To change the name, overwrite the existing name and click *Save* in the toolbar (see "Management Client Overview" on page 30). **Tip:** If you change the name, it will be updated throughout the system. This means that if the name is used in, for example, a rule, the name will automatically change in the rule as well.

Name	Description
Description	Description of the item. Optional. Will appear in a number of listings within the system. For example, the description will appear when pausing the mouse pointer over the item's name in the Overview pane (see "Panes Overview" on page 33):  Example from a camera... To specify a description, type the description and click ***Save*** in the toolbar (see "Management Client Overview" on page 30).
Hardware name	(only relevant for camera, microphone and speaker) Name of the hardware, with which the item is connected. The field is non-editable from here, but can be changed by clicking *Go To* next to it. This takes you to hardware information, where the name is editable.
Unit number	(only relevant for camera, microphone and speaker) Non-editable field displaying the unit on which the item is attached on the hardware. For single-device hardware, the unit number will typically be *1*. For multi-device hardware, such as video servers with several channels, the unit number will typically indicate the channel on which the item is attached, e.g. *3*.
Shortcut	This field is not in use.
Serial number	(only relevant for hardware) Hardware serial number as specified by the manufacturer. The serial number is often, but not always, identical to the MAC address.
Version	(only relevant for hardware and video encoders) Firmware version of the system as specified by the manufacturer. For a video en coder, it is the firmware version of the remote site system.
MAC address	(only relevant for hardware and video encoders) Hardware Media Access Control (MAC) address of the system hardware. A MAC address is a 12-character hexadecimal number uniquely identifying each device on a network. For a video encoder, it is the MAC address of the remote site system hardware.
Model	(only relevant for hardware and video encoders) Identifies the hardware model. For a video en coder, it identifies which remote site product the video encoder is communicating with.
Driver	(only relevant for hardware and video encoders) Identifies the driver handling the connection to the hardware. For a video en coder, it is the driver handling the connection to the remote site hardware.
Software license code	(only relevant for video encoders) Software license code of the of the remote system.
Windows user name	(only relevant for video encoders) Enter Windows user name for access to the remote site.
Windows password	(only relevant for video encoders) Enter Windows password for access to the remote site.

Name	Description
Connect	(only relevant for video encoders) When clicked, this opens a remote connection to the remote site (if Windows credentials are approved).

Settings tab overview

If you select a device group with 400 or more items the *Settings* tab will not be available for editing because changing settings for so many devices in one go takes too long time.

The content of the *Settings* tab is displayed in a table, in which the first column lists the available settings, and the second column lists the value of each setting. Values may be changeable or read-only. When you have changed a setting to a non-default value, the value will appear in **bold**. When a value must be within a certain range, for example between 0 and 100, the allowed range will be displayed in the information box below the settings table:

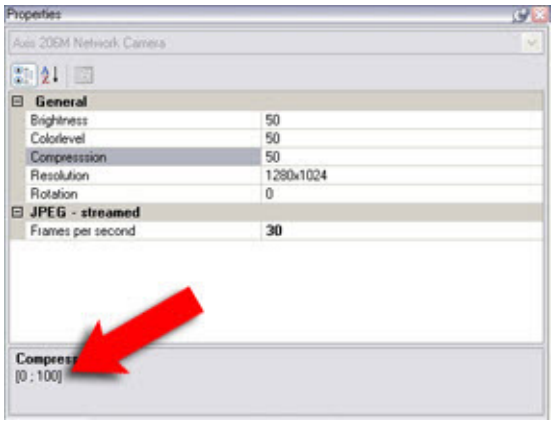

Settings tab, example from **camera**. Red arrow indicates allowed range; in this example the value used to specify compression must be a number between 0 and 100. Content of *Settings* tab varies depending on selected device type and selected device.

Tip: Some organizations may be required to establish a secure HTTPS connection using SSL (Secure Sockets Layer). To establish such a connection, you must upload a certificate to the hardware device to enable HTTPS support on the hardware device. Certificates are generated differently by camera vendors. Consult your camera vendor to find out how to get a certificate for your hardware device.

1. In the Management Client's Overview pane (see "Panes Overview" on page 33), right-click the required recording server to see its device groups. Select the relevant hardware under the wanted device group. On the *Settings* tab, all settings which are common to all of the device group's hardware will be listed.

Selecting hardware under a recording server

2. Select if you want to enable HTTPS on the hardware device. This is not enabled by default.
3. Enter the port to which the HTTPS connection is connected. The port number can be any numeric value between 1 and 65535.
4. Make changes as needed

5. Click *Save*.

HTTPS is enabled for the entire hardware device, that is, for example, a hardware device's camera, microphone and speaker.

CAMERA

Lets you view or edit settings, such as default frame rate, resolution, compression, the maximum number of frames between keyframes, on-screen date/time/text display, etc., for a selected camera, or for all cameras within a selected device group.

The content of the *Settings* tab is determined entirely by the drivers for the cameras in question, and is likely to vary depending on the types of cameras selected.

Some cameras may support more than one type of stream, for example MPEG4 and MJPEG. In that case, you can use multi-streaming (see "Streams tab (camera properties)" on page 101). If you change a camera's settings, you can quickly verify the effect of your change if you have the Preview pane (see "Panes Overview" on page 33) enabled. Note, however, that you cannot use the Preview pane to judge the effect of frame rate changes, as a special frame rate for the Preview pane's thumbnail images is used (defined in the Options dialog (see "Options" on page 207)).

Changing the settings for **Max. frames between keyframes** and **Max. frames between keyframes mode** may lower performance in the Ocularis Client.

MICROPHONE AND SPEAKER

Lets you verify or edit settings for selected microphones and speakers, or for all microphones or speakers within a selected device group.

Content of the *Settings* tab may vary depending on the types of microphones or speakers selected.

HARDWARE

Lets you verify or edit settings for the hardware selected under a recording server.

The content of the *Settings* tab is determined entirely by the hardware in question, and may vary depending on the type of hardware selected. For some types of hardware, the *Settings* tab may display no content at all or read-only content.

SPECIFY COMMON SETTINGS FOR ALL ITEMS IN A DEVICE GROUP—CAMERAS, MICROPHONES AND SPEAKERS

If using Device Groups (see "About device groups" on page 104), you are able to quickly specify common settings for all devices within a given device group:

1. In the list of device in the Management Client's Overview pane (see "Panes Overview" on page 33), right-click the required device group. On the *Settings* tab, all settings which are common to all of the device group's items (i.e. cameras, microphones or speakers) will be listed.
2. You are now able to verify or change both common settings and settings for individual item types within the device group.

Example are from camera.

From the menu above the settings list, select the required type of item:

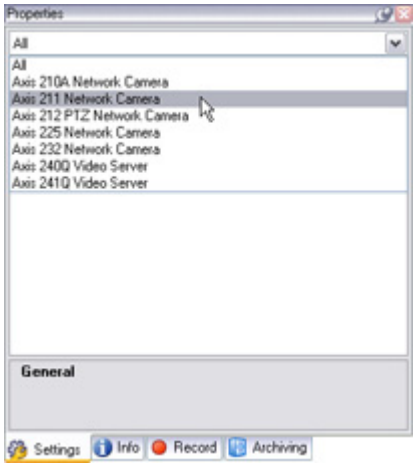

Example are from camera.

3. Make changes as needed.

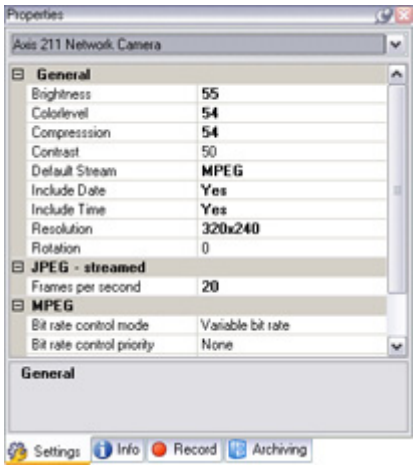

Example are from camera.

4. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

SPECIFY COMMON SETTINGS FOR ALL ITEMS IN A DEVICE GROUP—HARDWARE

1. In the Management Client's Overview pane (see "Panes Overview" on page 33), right-click the required recording server to see its device groups. Select the relevant hardware under the wanted device group. On the *Settings* tab, all settings which are common to all of the device group's hardware will be listed.

Selecting hardware under a recording server

2. You are now able to verify or change both common settings and settings for the individual hardware types within the device group.

   From the menu above the settings list, select the required type of hardware.

3. Make changes as needed
4. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

SET UP A SECURE CONNECTION ON ALL ITEMS IN A DEVICE GROUP

Tip: Some organizations may be required to establish a secure HTTPS connection using SSL (Secure Sockets Layer). To establish such a connection, you must upload a certificate to the hardware device to enable HTTPS support on the hardware device. Certificates are generated differently by camera vendors. Consult your camera vendor to find out how to get a certificate for your hardware device.

1. In the Management Client's Overview pane (see "Panes Overview" on page 33), right-click the required recording server to see its device groups. Select the relevant hardware under the wanted device group. On the *Settings* tab, all settings which are common to all of the device group's hardware will be listed.

Selecting hardware under a recording server

2. Select if you want to enable HTTPS on the hardware device. This is not enabled by default.
3. Enter the port to which the HTTPS connection is connected. The port number can be any numeric value between 1 and 65535.
4. Make changes as needed
5. Click *Save*.

HTTPS is enabled for the entire hardware device, that is, for example, a hardware device's camera, microphone and speaker.

Record tab overview

Recordings from an item (camera, microphone, speaker or OnSSI Interconnect remote systems (see "About OnSSI Interconnect" on page 57) will only be saved in the item's database when recording is enabled and recording-related rule (see "Manage rules" on page 165) criteria are met.

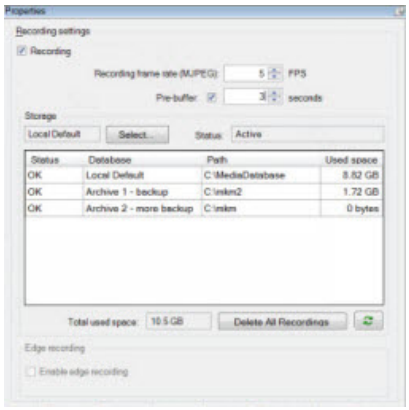

Record tab, example from **camera**

CAMERA

Lets you specify recording and storage settings for the selected camera.

What does recording mean? In IP video surveillance systems, the term *recording* means *saving video from a camera in the camera's database on the surveillance system*. In many IP video surveillance systems, all of the video received from cameras is not necessarily saved. Instead, saving of video in a camera's database, i.e. recording, is started only when there is a reason to do so: For example when motion is detected, when an event occurs, or when a specific period of time begins. Recording is then stopped after a specified amount of time, when motion is no longer detected, when an event occurs, when a time period ends, or similar. The term *recording* originates from the analog video era, when video was taped only when the record button was pressed.

MICROPHONE

Lets you specify recording and storage settings for the selected microphone. Microphones' recording and storage settings are completely independent of cameras and speakers.

SPEAKER

Lets you specify recording and storage settings for the selected speaker.

PLAYBACK - REMOTE SYSTEM

Playback settings are only visible if the selected hardware is part of an OnSSI Interconnect remote system (see "About OnSSI Interconnect" on page 57).

To enable viewing of playback directly from remote sites, select ***Play back recordings from remote system*** (see "***Enable playback directly from remote site camera***" on page 60).

Selecting this option disables the rest of the ***Recording settings*** options on the ***Record*** tab (see "Record tab overview" on page 125).

RECORDING

Recording is by default enabled. To enable/disable recording for the selected item, select/clear the *Record* tab's *Recording* check box.

Recording must be enabled for the item before you are able to record (i.e. save) video or audio from the camera. A rule (see "Manage rules" on page 165) specifying that an item should record under particular circumstances will not work if recording is disabled for the item in question.

RECORDING FRAME RATE - CAMERA

Specifying recording frame rate is only possible for MJPEG, a video codec (technology for compressing and decompressing data) with which each frame is separately compressed into a JPEG image.

1. Select or type the required recording frame rate (in FPS, Frames Per Second) in the ***Recording frame rate (MJPEG)*** box.
2. Clicking the ***Recording frame rate (MJPEG)*** box' up/down arrows will increase/reduce the value in increments of 1 FPS.

Tip: If you click inside the *Recording frame rate* box, two decimals will be added to the value. By selecting the number before or after the separator, you are able to increase/reduce the numbers in increments of 1 unit. This way you are able to specify a very specific recording frame rate average over time, for example of 20.15 FPS:

Specifying a specific recording frame rate

PREBUFFER

Prebuffering is essentially the ability to save video and audio in the camera's or microphone's database before the initial boundaries of a recording.

It can be highly advantageous as it allows you to save video and audio from **before** the events or times used to start recordings.

Cameras and microphones:

If, for example, you have created a rule specifying that recording should start when a door is opened, being able to see what happened immediately prior to the door being opened may be useful. Such prebuffering is possible since the system continuously receives streams of video and audio from connected cameras and microphones (unless the transfer of video or audio from cameras or microphones has in some way been disabled). Storing video and audio from before the initial boundaries of a recording is therefore not a problem: video and audio passes through the system anyway.

When prebuffering is enabled for a camera or a microphone, the system continuously records video or audio from the camera's or microphones stream and provisionally stores it in the database for a specified number of seconds before automatically deleting it— unless the provisionally stored video or audio turns out to be required for a recording, in which case it is automatically added to the recording.

Speakers:

Unlike video and incoming audio, which the system continuously receives from connected cameras and microphones, outgoing audio is only transmitted when Ocularis Client users press a button to talk through speakers. This can, depending on which events or times are used to start recordings, mean that there will be very little or no outgoing audio available for prebuffering.

Illustration of how prebuffered video/audio is added to a recording:

This is the stream received by the system:

These are the initial boundaries of a recording, as defined, for example, by start and stop events:

However, a rule specifies that recording should start 5 seconds prior to the start event, so 5 seconds of prebuffered video or audio is added:

This is what is actually recorded:

Enable and disable prebuffering:

Prebuffering is by default enabled; with a prebuffer size of 3 seconds. To enable/disable prebuffering, select/clear the *Pre-buffer **(in seconds)*** check box. When enabling, remember to specify a prebuffer size.

Specify prebuffer:

Select or type the required prebuffer size (in seconds) in the *Pre-buffer check* box. Clicking the *Pre-buffer* box' up/down arrows will increase/reduce the value in increments of one second.

The number of seconds you specify in the *Pre-buffer* checkbox must be sufficiently large to accommodate your requirements.

> **Example:** If, like in this rule example, you plan to be able to save video from five seconds prior to detected motion, the prebuffer size must be at least five seconds.

Use prebuffer in rules:

The use of prebuffering enables you to create rules (see "Manage rules" on page 165) specifying that recording should begin prior to the event or time triggering the rule.

> **Example:** Your ability to use this example rule— specifying that recording should start on a camera 5 seconds before motion is detected on the camera— depends on prebuffering being enabled for the camera in question.

Detail from a rule relying on prebuffering

STORAGE AREA

In the ***Storage*** area, you can monitor and edit database settings for the selected item.

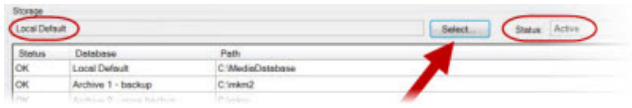

At the top of the *Storage* area, the selected database for the item in question and its status is stated. In this example, the selected database is *Local Default* and its status is *Active*.

Possible statuses for selected database:

Name	Description
Active	Database is active and running.
Archives also located in old storage	Database is active and running, and has archives located in other storage areas as well.
Data for some of the devices chosen is currently moving to another location	Database is active and running and moving data from one or more selected devices from one location to another.
Data for the device is currently moving to another location	Database is active and running and moving data from the selected device is currently moving from one location to another.
Information unavailable in failover mode	Status information about the database cannot be collected when database is in failover mode. For more information, see Manage failover recording servers (see "About failover recording servers—regular and hot standby" on page 234).

Further down in the *Storage* area, you can see which archive(s) are associated with the selected database, their individual status (***OK*** or ***Old Storage***), location and how much space they each use.

In the *Total used space* field, the total spaced used for the entire storage is indicated.

REMOTE RECORDING - CAMERA/REMOTE SYSTEM

The remote recording option is only available if the selected camera supports remote storage or is a camera under an OnSSI Interconnect remote site.

What is remote recording? *Remote recording* (also known as edge recording) is both a physical camera supporting edge storage and a remote recording system in an OnSSI Interconnect setup. To minimize loss if a network breaks down, some physical cameras are able to store recordings on their own local storage. Either on request or automatically (depending on settings), recordings can be retrieved from remote storages to the surveillance system when the network is re-established. To save bandwidth it is possible to set up rules regarding when to retrieves recordings.
With **remote systems,** the principle is the same. However, recordings are **not** retrieved from remote cameras' edge storages, but from remote systems' recording servers.

Select *Automatically retrieve remote recordings when connections are restored* (see "*Retrieve remote recordings from remote site camera*" on page 60) to enable automatic retrieval of recordings once connection is re-established.

The type of hardware selected determines where recordings are retrieved from:

- For a camera with local recording storage, recording are retrieved from the camera's local recording storage.
- For an OnSSI Interconnect remote system, they are retrieved from the remote systems' recording servers (see "About OnSSI Interconnect" on page 57).

Note that the *Retrieve remote recordings from <devices>* rule (see "Actions and Stop actions" on page 136) can be used independently of this setting.

Events tab overview

On the *Events* tab, you can manage events from hardware, cameras, microphones or OnSSI Interconnect remote systems (see "About OnSSI Interconnect" on page 57).

1. In the Overview pane (see "Panes Overview" on page 33) select the required item, in the Properties pane (see "Panes Overview" on page 33) select the *Events* tab.

   For hardware, the Events *tab will only be available if the selected item supports events.*

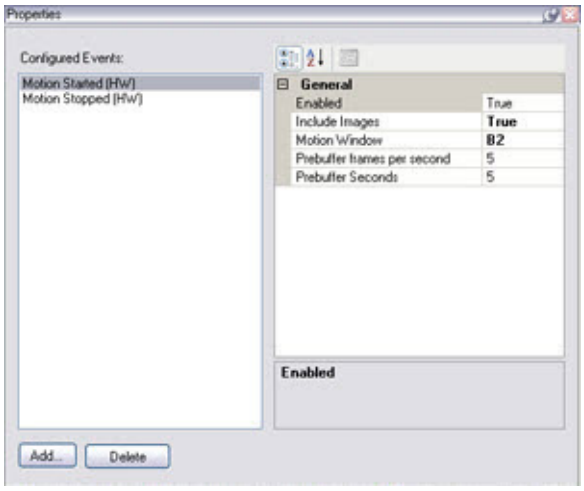

Event tab, example from **camera**

2. In the ***Configured Events*** list, add the wanted event(s) on each item on your system.

Which events you may select and add in the ***Configured events*** list is determined entirely by the hardware/device in question and its configuration. For some types of hardware/devices, the list may be empty. In OnSSI Interconnect setups, some events originating from remote systems will be predefined in the event list, but user-defined (see "Manage user-defined events" on page 180) events must be added manually. To view events added after an OnSSI Interconnect setup is established, you must update your remote site hardware (see "Update remote site hardware" on page 60).

CAMERA

In addition to the system's motion detection, some cameras can themselves be configured to detect motion. If a camera is capable of such detection, the camera's detections can be used as events. These events can be used when creating event-based rules (see "Events overview" on page 161) in the system. Technically, they occur on the actual hardware/device rather than on the surveillance system.

Events based on signals from input and/or output units connected to camera devices are managed elsewhere. Refer to Manage inputs (see "Manage input" on page 109) and Manage outputs (see "Manage output" on page 113).

In OnSSI Interconnect setups, some events originating from remote systems will be predefined in the event list, but user-defined (see "Manage user-defined events" on page 180) events must be added manually. To view events added after an OnSSI Interconnect setup is established, you must update your remote site hardware (see "Update remote site hardware" on page 60).

MICROPHONE

Some microphones are capable of creating events themselves. These events can be used when creating event-based rules (see "Events overview" on page 161) in the system. Technically, they occur on the actual hardware/device rather than on the surveillance system.

HARDWARE

Some hardware is capable of creating events itself. These events can be used when creating event-based rules (see "Events overview" on page 161) in the system. Technically, they occur on the actual hardware/device rather than on the surveillance system.

In OnSSI Interconnect setups, some events originating from remote systems will be predefined in the event list, but user-defined (see "Manage user-defined events" on page 180) events must be added manually. To view events added after an OnSSI Interconnect setup is established, you must update your remote site hardware (see "Update remote site hardware" on page 60).

ADD AN EVENT

1. On the *Events* Tab, click *Add...*. This opens the *Select Driver Event* window.
2. Select the required event. You can only select one event at a time.
3. Click *OK*. The selected event will be added to the *Events* tab's list of configured events.
4. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

Note that deleting an event (when possible) affects any rules in which the event is used.

In OnSSI Interconnect setups, some events originating from remote systems will be predefined in the event list, but user-defined (see "Manage user-defined events" on page 180) events must be added manually. To view events added after an OnSSI Interconnect setup is established, you must update your remote site hardware (see "Update remote site hardware" on page 60).

USE SEVERAL INSTANCES OF AN EVENT

To be able to specify different properties for different instances of an event (see "Specify event properties" on page 130), you are able to add an event more than once.

The following example is specific to **cameras**.

Example: The camera in question has been configured with two motion windows, called A1, and A2. You have added two instances of the *Motion Started (HW)* event. In the properties of one instance, you have specified use of motion window A1; in the properties of the other instance, you have specified use of motion window A2.

When you use the event in a rule, you are able to specify that the event should be based on motion detected in a specific motion window in order for the rule to be triggered:

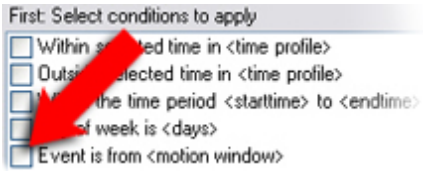

Example: Specifying specific motion window as part of a rule's conditions

SPECIFY EVENT PROPERTIES

For each event you have added, you are able to specify properties. The number of properties depends on the item in question. In order to work as intended, some or all of the properties must be specified identically on the item as well as on the system.

Even though the following list is not exhaustive, you may often be able to specify the following properties:

Name	Description
Enabled	Determines whether use of the event is enabled. Select *True* to enable; select *False* to disable. *Enabled* is the only property you will always see for microphones.
Include Images	Determines whether video should be sent from the camera to the system when the event occurs. Select *True* if video is required; select *False* if video is not required.

Name	Description
Motion Window	Many cameras capable of detecting motion can be configured with different motion detection settings for different parts the camera's images. For example, if a camera covers a 2-lane road, different motion detection settings may have been defined for the right lane and left lane area of the camera's images. Such areas are generally known as motion windows. Provided one or more motion windows have been defined on the camera device, the *Motion Window* setting lets you specify which motion window to use for the event. When the camera detects motion within the specified motion window, the event will occur. When specifying use of a motion window, make sure you type the name of the motion window, exactly as it has been specified on the camera. You can only specify one motion window in the field. However, you are able to add more than one instance of an event (see "Use several instances of an event" on page 130).
Prebuffer frames per second	Determines the frame rate to be used for prebuffered video. See also the next description of *Prebuffered Seconds* setting.
Prebuffer Seconds	Determines the number of seconds for which video from the camera should be stored for possible later use.

What does prebuffer mean? Prebuffering is essentially the ability to store video from before the initial boundaries of a recording. It allows you to view video from *before* an event occurred. If, for example, you are going to use the event in an rule specifying that recording should start when the event occurs, being able to see what happened immediately prior to the door being opened may also be important. An example could be, if you are using five seconds of prebuffering, video from the camera will always be stored provisionally for five seconds. If the event occurs, five seconds' worth of video will be available for attaching to any recording triggered by the event, as specified in a rule.

PTZ tab (video encoders)

The *PTZ* tab lets your enable PTZ (Pan/Tilt/Zoom) for video encoders. It is only available if the selected hardware is a video encoder.

What is a video encoder? A video encoder, also known as video server, is a piece of hardware which is able to stream video from a number of connected cameras. Video encoders contain image digitizers, making it possible to connect analog cameras to a network.

For video encoders, the use of PTZ must be enabled on the hardware level before you can use the PTZ features of PTZ cameras attached to the video encoder. The *Settings* tab lets you enable the use of PTZ separately for each of the video encoder's channels.

To access the *PTZ* tab, select the required hardware in the Overview pane (see "Panes Overview" on page 33), then select the *PTZ* tab in the Properties pane (see "Panes Overview" on page 33).

Not all video encoders support the use of PTZ cameras. Even video encoders which support the use of PTZ cameras may require configuration, such as installation of additional drivers (typically through accessing a browser-based configuration interface on the device's IP address) before PTZ cameras can be used.

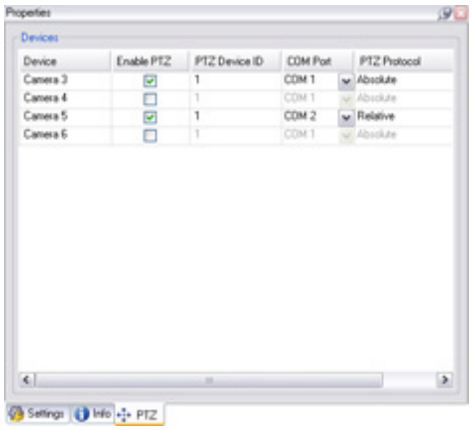

PTZ tab, with PTZ enabled for two of a video encoder's channels

ENABLE PTZ ON A VIDEO ENCODER

To enable the use of PTZ cameras on a video encoder, do the following on the *PTZ* tab:

1. In the list of devices connected to the video encoder, select the *Enable PTZ* box for the camera(s) on which you want to use PTZ:

2. In the *PTZ Device ID* column, verify the ID of the PTZ camera(s) in question.
3. In the *COM Port* column, select which of the video encoder's COM (serial communications) ports should be used for controlling PTZ functionality on each required PTZ camera:

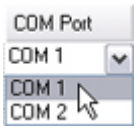

4. In the *PTZ Protocol* column, select which positioning scheme to use for each required PTZ camera:

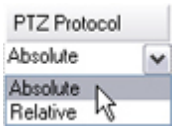

   - ***Absolute*** **:** When operators use Pan/Tilt/Zoom controls for the camera, the camera is adjusted relative to a fixed position, often referred to as the camera's home position
   - ***Relative*** **:** When operators use Pan/Tilt/Zoom controls for the camera, the camera is adjusted relative to its current position

   Refer to the camera's documentation if in doubt.

5. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

   You are now able to configure preset positions (see "PTZ Presets tab (camera properties)" on page 96) and patrolling (see "PTZ Patrolling tab (camera properties)" on page 93) for the PTZ camera(s) in question.

Remote Retrieval tab

The ***Remote Retrieval*** tab lets you handle remote recording retrieval settings for remote site hardware in an OnSSI Interconnect setup (see "About OnSSI Interconnect" on page 57):

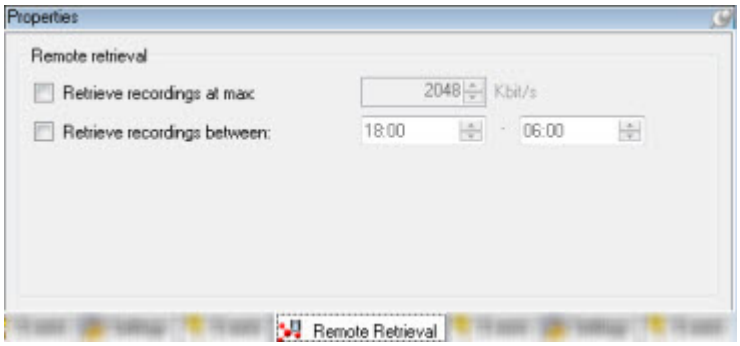

Specify the following properties:

- ***Retrieve recordings at max:*** Determines the maximum bandwidth in Kbits/s to be used for retrieving recordings from a remote site camera. Select the check box to enable limiting retrievals.
- ***Retrieve recordings between:*** Determines that retrieval of recordings from a remote site camera should be limited to a specific time interval.

None of the above applies to direct playback of remote recordings.

Note that if an automatic retrieval—or request for retrieval from the Ocularis Client—is received outside the time interval specified on the ***Remote Retrieval*** tab, it will be accepted, but not started until the selected time interval is reached. New remote recording retrieval jobs will queue and start when the allowed time interval is reached. Pending remote recording retrieval jobs can be viewed from the System Dashboard's Current Tasks (see "About current task" on page 195).

Status icons overview

The following icons are used to indicate status of cameras (see "Manage cameras" on page 82), microphones (see "Manage Microphones" on page 102), speakers (see "Manage speakers" on page 107), input (see "Manage input" on page 109) and output (see "Manage output" on page 113) events in item lists:

Cam-era	Micro phone	Spea-ker	In-put	Out-put	Description
					**Item enabled:** Can communicate with the recording server, and can if required be started/stopped automatically through a rule.
					**Item recording.** **Speaker being recorded:** Note that what is being said through the speaker can be recorded, but cannot subsequently be played back or exported (for example to prove that a warning was given).
					**Item temporarily stopped or has no feed:** Often shown when an item is communicating with the system while it is being disabled or enabled. Also shown if the Default Start Audio Feed Rule is not active; see Managing Rules. When stopped, no information is transferred to the system. In which case—if it is a camera—neither live viewing nor recording will be possible. However, a stopped item will still be able to communicate with the recording server, as opposed to when an item is disabled.

Cam-era	Micro phone	Spea-ker	In-put	Out-put	Description
					**Item disabled:** Cannot be started automatically through a rule and will not be able to communicate with the recording server. In the case of a camera, when a camera is disabled, neither live viewing nor recording will be possible.
					**Item database being repaired.**
					**Item requires attention.**
					**Status unknown.**
					Note that some icons may be combined, as in this example where **Item is enabled** is combined with **Item is recording** (since a recording item is also an enabled item).

Clients

About clients

In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), you are able to work with the following under *Clients*:

- (see "Manage view groups" on page 135) **:** Here you manage your View Groups, which are basically containers for one or more logical groups of views.

About view groups

The way in which video from one or more cameras is presented in clients is called a view. A view group is basically a container for one or more logical groups of such views.

Manage view groups

By default, each role you define in the Management Client is also created as a view group. When you add a role in the Management Client, the role will by default appear as a view group for use in clients.

- A view group based on a role will by default only be available to users/groups assigned to the role in question. You may change these view group rights (on page 192).
- A view group based on a role will by default carry the role's name.

  In addition to the view groups you get when adding roles, you may create as many other view groups as you like. You can also delete view groups, including those automatically created when adding roles.
- Even though a view group is created by default each time you add a role (see "Manage roles" on page 184), view groups do not have to correspond to roles. You may therefore add, rename or remove any of your view groups if required.

Note that if you rename a View group, client users already connected must log out and log in again before the name change will be visible.

ADD A VIEW GROUP

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand the ***Clients*** node, right-click ***View Groups***, and select *Add View Group*. This opens the *Add View Group* dialog.
2. Type the name of the new view group, then click *OK*.
3. Optionally; in the Management Client's Overview pane (see "Panes Overview" on page 33), select the added view group, then in the Properties pane (see "Panes Overview" on page 33) add a description of the view group.

No roles will have the right to use the newly added view group until you have specified such rights; refer to View group rights (on page 192) for more information.

Also, even when you have specified which roles should be able to use the newly added view group, already connected client users with the relevant roles must log out and log in again before they will be able to see the view group.

Rules and events

About rules and events

In your system, events are central elements when using the *Manage Rule* wizard (see "Manage rules" on page 165). In the wizard, events are primarily used for triggering actions**.**

> **Example:** You create a rule which specifies that in the *event* of detected motion, the surveillance system should take the *action* of starting recording of video from a particular camera.

In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), you can work with the following under *Rules and Events*:

- ***Rules*** (see "***Manage rules***" on page 165) ***:*** Rules are a central element in the system. The behavior of your surveillance system is to a very large extent determined by rules. When creating a rule, you can work with all types of events.
- ***Time profiles*** (see "***Manage time profiles***" on page 173) ***:*** Time profiles are periods of time defined in the Management Client. They can be used when creating rules in the Management Client; for example, to create a rule which specifies that a certain action should take place within a certain time profile.
- ***Notification profiles*** (see "***Manage notification profiles***" on page 176) ***:*** Notification profiles can be used for setting up ready-made e-mail notifications, which can automatically be triggered by a rule, for example when a particular event occurs.
- ***User-defined events*** (see "***Manage user-defined events***" on page 180) ***:*** User-defined events are custom-made events making it possible for users to manually trigger events in the system or react to inputs from the system.

Actions and Stop actions

Depending on the recording component, functionality described here may be limited or unavailable.

When using events to create rules in the *Manage Rule* wizard (see "Manage rules" on page 165), you are able to select between a number of different actions:

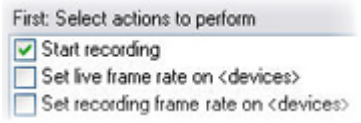

Example: Selecting actions

Some of these actions will require a subsequent stop action.

> ***Example:*** If you select the action *Start recording*, recording will start and potentially continue indefinitely. Therefore, the action *Start recording* has a compulsory stop action called *Stop recording*.

Manage Rule makes sure you specify such stop actions when necessary; stop actions are typically specified on one of the last steps of the wizard:

Selecting stop actions. In the example, note the compulsory stop action (selected, dimmed), the non-relevant stop actions (dimmed) and the optional stop actions (selectable).

Each type of action is described (additional actions may, however, be available if your system installation uses add-on products, special plug-ins, etc.). For each type of action, stop action information is listed as well:

Action	Description
Start recording	Begin recording, i.e. begin saving video in the database of the selected camera. When selecting this type of action, *Manage Rule* will prompt you to specify when recording should start (either immediately or a number of seconds before the triggering event/beginning of the triggering time interval) as well as on which devices the action should take place. This type of action requires that recording has been enabled on the cameras to which the action will be linked. Being able to save video from before an event or time interval is only possible if prebuffering is enabled for the cameras in question. You enable recording and specify prebuffering settings for a camera on the ***Record*** tab (see "Record tab overview" on page 125). ***Stop action required:*** This type of action requires one or more stop actions. Depending on how the action was triggered, the stop action may be performed either on an event or after a period of time. In one of the subsequent steps of Manage Rule, the wizard will automatically prompt you to specify the stop action: Stop recording. Without this stop action, recording would potentially continue indefinitely. You also have the option of specifying further stop actions.
Start feed on <devices>	Begin video feed from camera devices to the system. When the feed from a device is started, video will be transferred from the device to the system , in which case live viewing and recording of video is possible. **IMPORTANT:** While this type of action enables access to selected cameras' video feeds, it does not guarantee that video is recorded, as cameras' recording settings must be specified separately. When selecting this type of action, the ***Manage Rule*** wizard (see "Manage rules" on page 165) wizard will prompt you to specify on which devices feeds should be started. **Tip:** Your system has a default rule ensuring that feeds are always started on all cameras. Note however, that the default rule may have been manually deactivated or modified. ***Stop action required:*** This type of action requires one or more stop actions. Depending on how the action was triggered, the stop action may be performed either on an event or after a period of time. In one of the subsequent steps of Manage Rule, the wizard will automatically prompt you to specify the stop action: Stop feed. You also have the option of specifying further stop actions. Note that using the compulsory stop action *Stop feed* to stop the feed from a device means that video will no longer be transferred from the device to the system, in which case live viewing and recording of video will no longer be possible. However, a device on which the feed has been stopped will still be able to communicate with the recording server, and the feed can be started again automatically through a rule, as opposed to when the device has been manually disabled in the Management Client.
Set live frame rate on <devices>	Sets a particular frame rate to be used when displaying live video from the selected cameras, instead of the cameras' default frame rate (specified on the ***Settings*** tab (see "Settings tab overview" on page 121)). When selecting this type of action, *Manage Rule* will prompt you to specify which frame rate to set, and on which devices. Always verify that the frame rate (number of frames per second) you specify is

Action	Description
	available on the cameras in question. ***Stop action required:*** This type of action requires one or more stop actions. Depending on how the action was triggered, the stop action may be performed either on an event or after a period of time. In one of the subsequent steps of Manage Rule, the wizard will automatically prompt you to specify the stop action: Restore default live frame rate. Without this stop action, the default frame rate would potentially never be restored. You will also have the option of specifying further stop actions.
Set recording frame rate on <devices>	Sets a particular frame rate to be used when saving recorded video from the selected cameras in the database, instead of the cameras' default recording frame rate. When selecting this type of action, *Manage Rule* will prompt you to specify which recording frame rate to set, and on which cameras. Specifying recording frame rate is only possible for MJPEG, a video codec (technology for compressing and decompressing data) with which each frame is separately compressed into a JPEG image. This type of action also requires that recording has been enabled on the cameras to which the action will be linked. You enable recording for a camera on the *Record* tab (see "Record tab overview" on page 125). The maximum frame rate you will be able to specify will depend on the camera types in question, and on their selected image resolution. ***Stop action required:*** This type of action requires one or more stop actions. Depending on how the action was triggered, the stop action may be performed either on an event or after a period of time. In one of the subsequent steps of Manage Rule, the wizard will automatically prompt you to specify the stop action: Restore default recording frame rate. Without this stop action, the default recording frame rate would potentially never be restored. You will also have the option of specifying further stop actions.
Start patrolling on <device> using <profile> with PTZ priority <priority>	Begins PTZ patrolling (the continuous moving of a camera between a number of preset positions) according to a particular patrolling profile (the exact definition of how patrolling should be carried out, including the sequence of preset positions, timing settings, etc.) for a particular PTZ camera with a particular priority. **What is Priority?** When several users on a surveillance system wish to control the same PTZ camera at the same time, conflicts may occur. PTZ priority lets you alleviate the problem by specifying a priority for use of the selected PTZ camera(s) by users/groups with the selected role. Specify a priority from *1* to *32.000*, where *1* is the lowest priority. Default PTZ priority is 3000. If your system is upgraded from an older version of the system , the old values (*Very Low, Low, Medium, High* and *Very High*) have been translated as follows: ○ Very Low = 1000 ○ Low = 2000 ○ Medium = 3000 ○ High = 4000 ○ Very High = 5000 If your system is upgraded to version 4.0 (or future versions), rule priority settings is a new feature. Existing rules (created without priority) automatically get priority 1. It is strongly recommended to reconsider this lowest possible priority for all affected rules. When selecting this type of action, *Manage Rule* will prompt you to select a patrolling profile. Only one patrolling profile on one device can be selected; it is not possible to select several patrolling profiles. This type of action requires that the device(s) to which the action will be linked is/are a

Action	Description
	PTZ (Pan/Tilt/Zoom) device. Furthermore, it requires that at least one patrolling profile has been defined for the device(s). You define patrolling profiles for a PTZ camera on the *Patrolling* tab (see "PTZ Patrolling tab (camera properties)" on page 93). ***Stop action required:*** This type of action requires one or more stop actions. Depending on how the action was triggered, the stop action may be performed either on an event or after a period of time. In one of the subsequent steps of Manage Rule, the wizard will automatically prompt you to specify the stop action: Stop patrolling. Without this stop action, patrolling would potentially never stop. You will also have the option of specifying further stop actions.
Pause patrolling on <devices>	Pauses PTZ patrolling (the continuous moving of a camera between a number of preset positions). When selecting this type of action, *Manage Rule* will prompt you to specify the devices on which patrolling should be paused. This type of action requires that the device(s) to which the action will be linked is/are a PTZ (Pan/Tilt/Zoom) device. Furthermore, it requires that at least one patrolling profile has been defined for the device(s). You define patrolling profiles for a PTZ camera on the *Patrolling* tab (see "PTZ Patrolling tab (camera properties)" on page 93). ***Stop action required:*** This type of action requires one or more stop actions. Depending on how the action was triggered, the stop action may be performed either on an event or after a period of time. In one of the subsequent steps of Manage Rule, the wizard will automatically prompt you to specify the stop action: Resume patrolling. Without this stop action, patrolling would potentially pause indefinitely. You will also have the option of specifying further stop actions.
Move <device> to <preset> position with PTZ priority <priority>	Moves a particular camera to a particular preset position - however always according to priority. When selecting this type of action, *Manage Rule* will prompt you to select a preset position. Only one preset position on one camera can be selected; it is not possible to select several preset positions. If your system is upgraded to version 4.0 (or future versions), rule priority settings is a new feature. Existing rules (created without priority) automatically get priority 1. It is strongly recommended to reconsider this lowest possible priority for all affected rules. This type of action requires that the device(s) to which the action will be linked is/are a PTZ (Pan/Tilt/Zoom) device. Furthermore, it requires that at least one preset position has been defined for those devices. You define preset positions for a PTZ camera on the *Presets* tab (see "PTZ Presets tab (camera properties)" on page 96). ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.
Move to default preset on <devices> with PTZ priority <priority>	Moves one or more particular cameras to their respective default preset positions - however always according to priority. When selecting this type of action, *Manage Rule* will prompt you to select which devices the action should apply for. If your system is upgraded to version 4.0 (or future versions), rule priority settings is a new feature. Existing rules (created without priority) automatically get priority 1. It is strongly recommended to reconsider this lowest possible priority for all affected rules. This type of action requires that the device(s) to which the action will be linked is/are a PTZ (Pan/Tilt/Zoom) device. Furthermore, it requires that at least one preset position has been defined for those

Action	Description
	devices. You define preset positions for a PTZ camera on the *Presets* tab (see "PTZ Presets tab (camera properties)" on page 96). ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.
Set device output to <state>	Sets an output on a device to a particular state (activated or deactivated). When selecting this type of action, *Manage Rule* will prompt you to specify which state to set, and on which devices. This type of action requires that the devices to which the action will be linked each have at least one external output unit connected to an output port. ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.
Create bookmark on <device>	Creates a bookmark on live streaming or recordings from a selected device. A bookmark makes it easy to retrace a certain event or period in time. Bookmark settings are controlled from the ***Options*** (on page 207) dialog. When selecting this type of action, *Manage Rule* will prompt you to specify bookmark details and select device. ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.
Send notification to <profile>	Sends a notification, using a particular notification profile. When selecting this type of action, *Manage Rule* will prompt you to select a notification profile, and which devices to include pre-alarm images from. Only one notification profile can be selected; it is not possible to select several notification profiles. **Tip:** Even though you are only able to select a single notification profile, bear in mind that a single notification profile may contain several recipients. This type of action requires that at least one notification profile (see "Manage notification profiles" on page 176) has been set up. Pre-alarm images are only included if e-mail notification is used and the *Include images* option has been enabled for the notification profile in question. ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.
Make new <log entry>	Generates an entry in the rule log (see "Manage logs" on page 197). When selecting this type of action, *Manage Rule* will prompt you to specify a text for the log entry. **Tip:** When specifying the log text, you will be able to quickly insert variables, such as **$DeviceName$, $EventName$**, etc. into the log message wording. ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.
Start plug-in on <devices>	Starts one or more plug-ins. When selecting this type of action, *Manage Rule* will prompt you to select required plug-ins, and on which devices to start the plug-ins. This type of action requires that at one or more plug-ins are available on your system. ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.
Stop plug-in on <devices>	Stops one or more plug-ins. When selecting this type of action, *Manage Rule* will prompt you to select required plug-ins, and on which devices to stop the plug-ins.

Action	Description
	This type of action requires that at one or more plug-ins are available on your system. ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.
Apply new settings on <devices>	Changes device settings. When you select this type of action, *Manage Rule* will prompt you to select required devices, and you will be able to define required settings on the devices you have specified. If defining settings for more than one device, you will only be able to change settings that are available for all of the specified devices. ***Example:*** You specify that the action should be linked to Device 1 and Device 2. Device 1 has the settings A, B and C, and Device 2 has the settings B, C and D. In this case, you will only be able to change the settings that are available for both devices, namely settings B and C. ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.
Set NetMatrix to view <devices>	Makes video from the selected cameras appear on a computer capable of displaying NetMatrix -triggered video. When you select this type of action, *Manage Rule* will prompt you to select a NetMatrix recipient, and one or more devices from which to display video on the selected NetMatrix recipient. This type of action lets you select only a single NetMatrix recipient at a time. If you want to make video from the selected devices appear on more than one NetMatrix recipient, you should create a rule for each required NetMatrix recipient. NetMatrix is not used with Ocularis.
Send SNMP trap	Generates a small message which logs events on selected devices. The text of SNMP traps are auto-generated and cannot be customized. It will typically contain the source type and name of the device on which the event occurred. To configure who receives SNMP trap messages, refer to SNMP support (see "About SNMP support" on page 251). ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.
Retrieve edge storage recordings from <devices>	Retrieves and stores edge recordings from selected devices (that support edge recording (see "Remote recording - camera/remote system" on page 128)). Can be set to execute immediately or at a certain point in time. ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time. Note that this rule is independent of the *Automatically retrieve remote recordings when connection is restored* setting (see "Remote recording - camera/remote system" on page 128).
Save attached image	Ensures that when an image is received from the Images Received event (see "Events overview" on page 161) (sent via SMTP e-mail from a camera) it is saved for future usage. In future, other events might also be able to trigger this action. ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.
Activate archiving on <archives>	Starts archiving on one or more archives. When you select this type of action, *Manage Rule* will prompt you to select required archives.

Action	Description
	No compulsory stop action: This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.
On <site> trigger <user-defined event>	Relevant mostly within OnSSI Federated Architecture (see "About OnSSI Federated Architecture" on page 212), but can also be used in a single server setup. Used for triggering a user defined event on a site - normally a remote site within a federated hierarchy. ***No compulsory stop action:*** This type of action does not require a stop action; although it will be possible to specify optional stop actions to be performed on either an event or after a period of time.

Create typical rules

The following is a brief introduction to examples of typical rules, what you can do with them, and how they can be constructed.

Basic rules

- ***Use higher live frame rate on motion*** *:* Ensures that when motion is detected on a specific camera, the system will use a higher than default live frame rate for the camera, and return to using the camera's default live frame rate when motion is no longer detected. The effect is higher quality live video whenever there is motion. When the specified part of the day ends, the PTZ camera will stop patrolling.

PTZ-related rules

- ***Use specific PTZ patrolling profile during specific part of day*** *:* Ensures that during a specific part of the day, a PTZ (Pan/Tilt/Zoom) camera will patrol according to a specific patrolling profile (i.e. the exact definition of how patrolling should be carried out, including the sequence for moving between preset positions, timing settings, etc.). When the specified part of the day ends, the PTZ camera will stop patrolling.
- ***Use different PTZ patrolling profiles for day/night:*** Ensures that during daytime, a PTZ camera will patrol according to a specific patrolling profile. And during nights, according to another patrolling profile.
- ***Pause PTZ patrolling and go to PTZ preset on input*** *:* Ensures that a specific external input is activated, a PTZ camera will pause its patrolling, move to a specific preset position, and remain at the preset position for a specific period of time, after which it will resume patrolling.

Use higher live frame rate on motion rule

In this example, the camera has a default live frame rate of 10 frames per second (FPS), and the rule increases the live frame rate to 25 FPS when applied. The effect is live video of a higher quality for as long as motion is detected on the camera.

Note that recording frame rate (the frame rate with which video sequences will be saved) is specified separately, and is not affected by this rule.

If you want to permanently change the default frame rate for a camera, do not use a rule. Change the camera's default frame rate on the **Settings** tab (see "Settings tab overview" on page 121) instead.

Motion is normally detected by the system when video received from cameras is analyzed. This is the type of motion detection dealt with in this example. However, some cameras are— depending on configuration of the camera hardware— themselves able to detect motion. Such motion detection can also be used in system rules, although that is beyond the scope of this example.

Prerequisites

This rule is based on motion detection on a specific camera. Therefore, motion detection must be enabled on the camera in order for the rule to work as intended. Before creating a rule like this, always verify the following:

- Motion detection is enabled for the camera in question

Show me how to verify this...

To verify that motion detection has been enabled for a camera, expand Devices in the Management Client's Site Navigation pane (see "Panes Overview" on page 33), and select ***Cameras***. This will display a list of cameras in the Overview pane (see "Panes Overview" on page 33). Select the required camera from the list, and select the ***Motion*** tab in the Properties pane (see "Panes Overview" on page 33). On the ***Motion*** tab, verify that the ***Enable motion detection*** check box is selected.

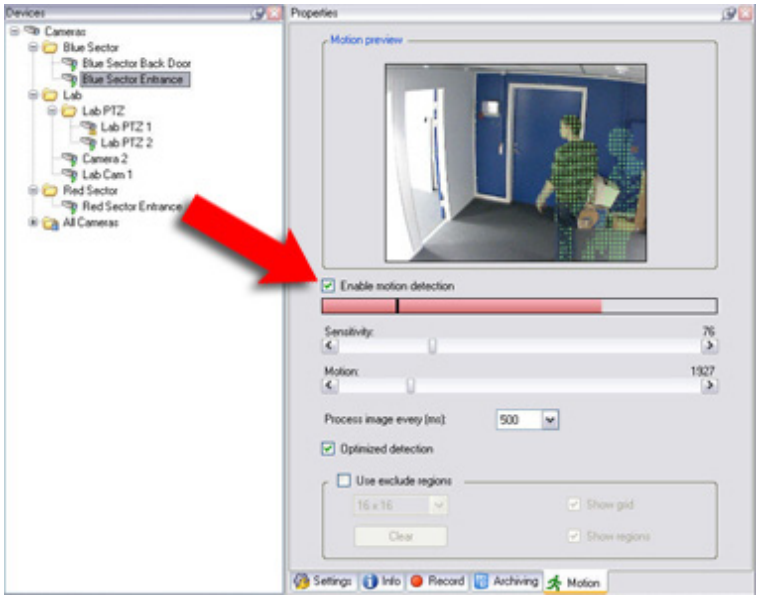

Arrow indicates position of *Enable motion detection* check box

Note that other settings on the *Motion* tab, such as *Sensitivity*, will determine what will be interpreted as motion. Merely enabling motion detection may thus not be sufficient to meet your requirements. Time spent on finding the best possible balance of motion detection settings under different conditions (day/night, calm/windy weather, etc.) will help you later avoid unnecessary recordings, etc.

Creating the Rule

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Rules and Events*, then right-click Rules and select *Add New Rule...*:

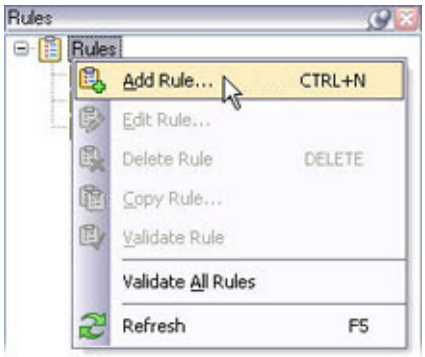

2. The *Manage Rule* wizard opens. Type a name for the new rule in the *Rule name* field.

   In this example... the rule will cover a specific camera, Camera 1. We therefore overwrite the default rule name (e.g. New Rule 001) with a descriptive name:

   

   Tip: Always use a descriptive name for the rule. Once you have several rules, you will find that descriptive names are a great help when identifying individual rules.

3. On Step 1 of *Manage Rule*, select the required rule type.

   In this example... we want to base the rule on an event, namely detected motion. Therefore, we select *Perform an action on <event>*. Our selection is immediately reflected in the initial rule description in the lower half of the wizard window:

   

4. Click the underlined items in the rule description in order to specify its exact content:

   Event link: Clicking the *event* link lets you select the event which must occur in order for the rule to apply. In order for you to get a good overview, selectable events are listed in groups according to whether they are related to plug-ins, dependent on hardware configuration or built into the system itself, etc.

   In this example ... we want the event to be detected motion. Motion detection events are technically related to the system's motion detection plug-in, so we go to the *Plug-ins* group, select the event *Motion Start*, and click *OK*:

   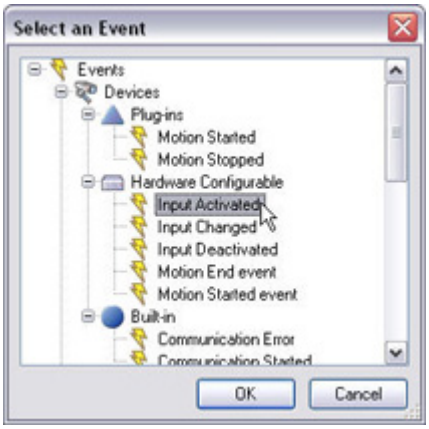

   Devices/recording server/management server link: When you have selected the required event, clicking the *devices/recording server/management server* link opens the *Select Group Members* window, which lets you specify the devices on which device the event should occur in order for the rule to apply.

   In this example ... the event should occur on Camera 1 in order for the rule to apply. In the *Select Group Members* window we therefore drag Camera 1 to the *Selected* list and click *OK*. By doing this we have specified the exact content of the first part of the wizard's rule description, which now looks like this:

   

5. Click *Next* to move to step 2 of the wizard. On step 2 of the wizard, specify which time conditions should be met in order for the rule to apply.

In this example ... we want the rule to apply whenever motion is detected on Camera 1, regardless of time. When creating event-based rules it is possible to bypass the time conditions; we therefore want to skip step 2 entirely.

6. Click *Next* to move to step 3 of the wizard. On step 3 of the wizard, first specify which actions to perform.

   In this example ... we want to set a specific live frame rate. We therefore select the action *Set live frame rate on <devices>*:

   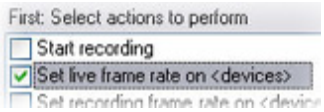

   Based on the selection of actions, the wizard automatically extends the rule description in the lower part of the wizard window.

   In this example ... Based on our selection *Set live frame rate on <devices>*, the wizard automatically suggests a rule description in which the frame rate should be set on *the device on which event occurred*. The wizard furthermore prompts us to specify the required number of frames per second:

   

   To specify the required number of frames per second, we click the *frames per second link*, specify a frame rate of 25, and click *OK*:

   

   The rule description now indicates that the frame rate will be set to 25 frames per second.

7. Click *Next* to move to step 4 of the wizard. On step 4 of the wizard, select stop criteria. Stop criteria are important in many types of rules. Without a stop criterion, many actions could go on indefinitely once started.

   In this example ... Without a stop criterion, the rule in this example would set the frame rate for the camera to 25 FPS indefinitely upon motion detection. Based on the elements in our rule description, the wizard therefore automatically suggests the stop criterion *Perform stop action on <event>*:

   

   Note that the stop criterion *No actions performed on rule end* is not available: a stop criterion must be defined for this type of rule. In the rule description, the wizard furthermore automatically suggests that the stop action is performed when motion is no longer detected on Camera 1:

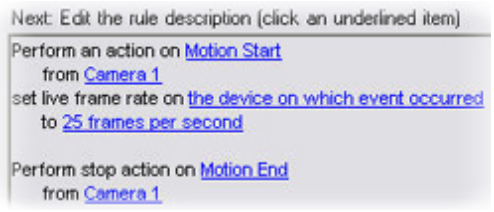

This is just what we want; we do not need to change any of the wizard's suggestions. However, we still need to define exactly which kind of stop action should take place when motion ends on Camera 1.

8. Click *Next* to move to the next step of the wizard. In this step, the wizard suggests one or more stop actions based on the previously selected start actions.

   In this example ... Based on the start action *set frame rate* in our rule description, the wizard automatically suggests the stop action *restore default frame rate*. It furthermore suggests that the default frame rate should be restored immediately after the last detected motion:

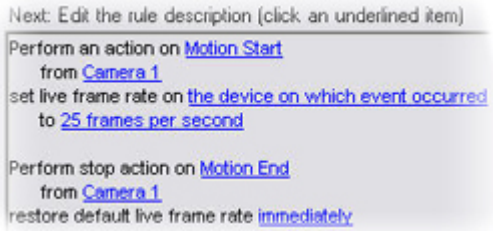

This is also just what we want; we do not need to change it, although by clicking the *immediately* link we could have specified a delay of e.g. 3 seconds.

Your rule is by default active, meaning that once you have created it, it will be applied as soon as the rule's conditions are met. If you do not want the rule to be active straight away, clear the *Active* check box in the top part of the *Manage Rule* window:

Tip: You can always activate/deactivate the rule later.

9. Click *Finish*. This will add your new rule to the list of rules.

Use specific PTZ patrolling profile during specific part of day rule

Tip: When patrolling stops, you can—if needed—get the PTZ camera to start patrolling immediately after according to another patrolling profile.

Prerequisites

When a PTZ camera patrols according to a patrolling profile, it continuously moves between different preset positions. Therefore, the required preset positions and at least one patrolling scheme must be defined for the PTZ camera in question. Before creating a rule like this, always verify the following:

- The camera in question is a PTZ camera
- At least two preset positions are defined for the camera

  How to define preset positions...

  To define preset positions for a PTZ camera, expand ***Devices*** in the Management Client's Site Navigation pane (see "Panes Overview" on page 33) and select ***Cameras***. In the Overview pane (see "Panes Overview"

on page 33), select the required PTZ camera from the list, then select the ***Presets*** tab in the Properties pane (see "Panes Overview" on page 33). For descriptions of the exact functionality of the ***Presets*** tab, refer to Preset positions (see "PTZ Presets tab (camera properties)" on page 96).

- At least one patrolling profile is defined for the camera

  How to define a patrolling profile...

  To define patrolling profiles for a PTZ camera, expand ***Devices*** in the Management Client's Site Navigation pane (see "Panes Overview" on page 33) and select ***Cameras***. In the Overview pane (see "Panes Overview" on page 33), select the required PTZ camera from the list, then select the ***Patrolling*** tab in the Properties pane (see "Panes Overview" on page 33). For descriptions of the exact functionality of the *Patrolling* tab, refer to Patrolling (see "PTZ Patrolling tab (camera properties)" on page 93).

Creating the Rule

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Rules and Events*, then right-click Rules and select *Add New Rule...*:

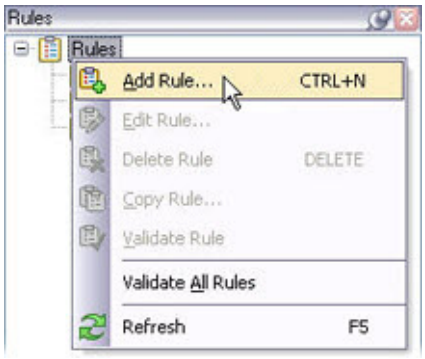

2. The *Manage Rule* wizard opens. Type a name for the new rule in the *Rule name* field.

   In this example... the rule will only cover a specific camera (called *PTZ Camera*) and how it should patrol on Saturday afternoons. We therefore overwrite the default rule name (e.g. *New Rule 001*) with a descriptive name:

   Tip: Always use a descriptive name for the rule. Once you have several rules, you will find that descriptive names are a great help when identifying individual rules.

3. On Step 1 of *Manage Rule*, select the required rule type.

   In this example... we want to base the rule on a time period. Therefore, we select *Perform an action in a time interval*:

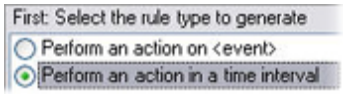

   Click *Next* to move to the next step of the wizard.

4. On the wizard's next step, specify which time conditions should be met in order for the rule to apply.

   In this example... we want the rule to apply between 1:00 and 8:00 on Saturdays, so two time conditions are required: one which specifies use of a start time and end time, and one which specifies use on a particular day of the week. We therefore select *Within the time period <start time> to <end time>* and *Day of week is <day>*:

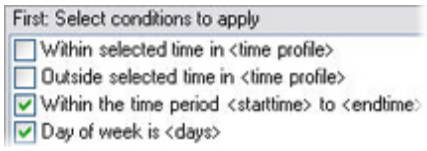

Our selection is immediately reflected in the initial rule description in the lower half of the wizard window:

Tip: If we had previously created a suitable time profile covering the required period of time, we could have just selected the time condition within selected time in *<time profile>*, then pointed to the time profile in question.

Read more about time profiles under Manage time profiles (on page 173).

5. Click the underlined items in the rule description in order to specify its exact content:

   start time: Clicking the start time link lets you specify required start time.

   In this example ... we want the start time to be one o'clock in the afternoon, so we specify 1:00, and click OK:

   end time: The *end time* link works just like the *start time* link. We specify 8:00.

   days: Clicking the *days* link lets you specify required days of the week.

   In this example ... our rule should only apply on Saturdays, so we select *Saturday*, and click *OK*:

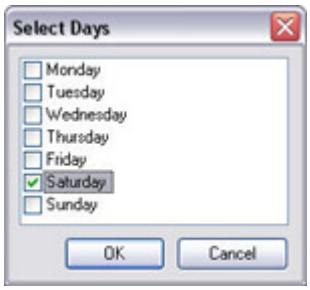

   By doing this, we have specified the exact content of the first part of the wizard's rule description, which now looks like this:

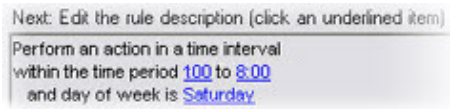

   Click *Next* to move to step 3 of the wizard.

6. On step 3 of the wizard, first specify which actions to perform.

In this example ... we want to start patrolling according to a specific patrolling profile. We therefore select the action *Start patrolling on <device> using <profile> with PTZ priority <priority>*:

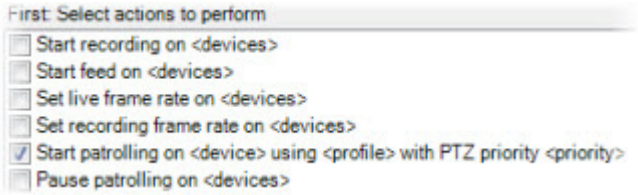

Based on the selection of actions, the wizard automatically extends the rule description in the lower part of the wizard window.

In this example... Based on our selection *Start patrolling on <device> using <profile> with PTZ priority <priority>*, the wizard automatically prompts us to specify the required camera, patrolling profile and its priority (see "Actions and Stop actions" on page 136):

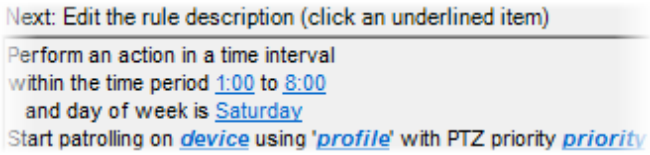

We click the *device* link, expand the relevant camera folder, select the required camera, and click *OK*:

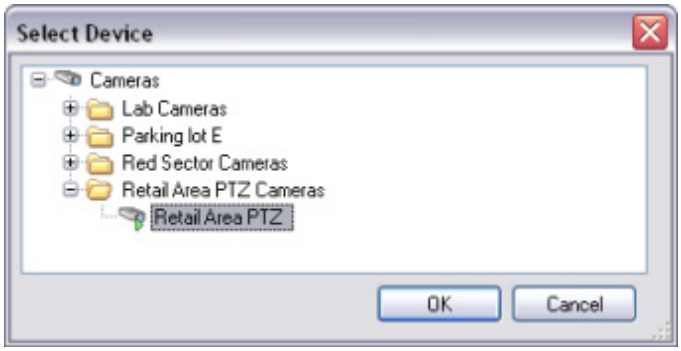

Next we click the *profile* link, select the required patrolling profile from our list of previously defined patrolling profiles, and click *OK*:

Finally, we click the *priority* link to set the priority (see "Actions and Stop actions" on page 136) of the patrolling profile.

By doing this, we have further specified the content of the wizard's rule description, which now looks like this:

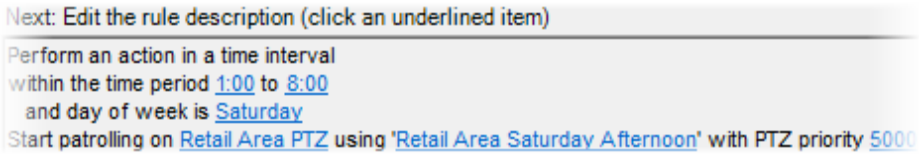

Click *Next* to move to step 4 of the wizard.

7. On step 4 of the wizard, select stop criteria. Stop criteria are important in many types of rules. Without a stop criterion, many actions could go on indefinitely once started.

   In this example... Without a stop criterion, the rule in this example would start patrolling within the specified time period, but never stop it. Based on the elements in our rule description, the wizard therefore automatically suggests the stop criterion *Perform stop action when time interval ends*:

   

   Note that the stop criterion *No actions performed on rule end* is not available: a stop criterion must be defined for this type of rule. We still need to define exactly which kind of stop action should take place when the time period ends.

   Click *Next* to move to the next step of the wizard.

8. In this step, the wizard suggests one or more stop actions based on the previously selected start actions.

   In this example ... Based on the start action start patrolling in our rule description, the wizard automatically suggests the stop action *Stop patrolling*. It furthermore suggests that patrolling is stopped immediately when the time period ends:

   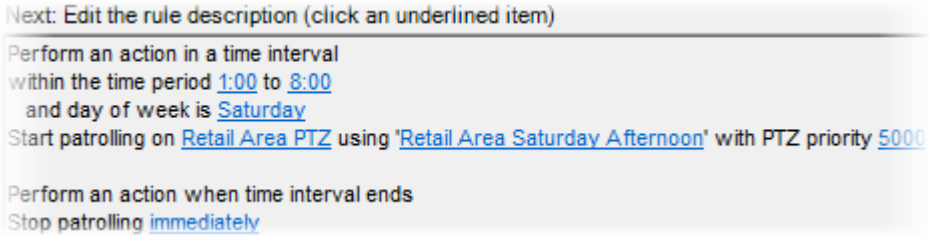

   This is just what we want; we do not need to change it, although by clicking the *immediately* link we could have specified a delay of e.g. 60 seconds.

   Your rule is by default active, meaning that once you have created it, it will be applied as soon as the rule's conditions are met. If you do not want the rule to be active straight away, clear the *Active* check box in the top part of the *Manage Rule* window:

   

   Tip: You can always activate/deactivate the rule later.

9. Click *Finish*. This will add your new rule to the list of rules:

   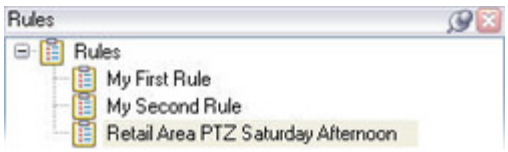

Use different PTZ patrolling profiles for day/night rule

In this example, *daytime* is defined by a time profile covering the period between 08.00 and 20.00 on all days of the week and *nights* are defined as periods not covered by the *daytime* time profile. This requires two near-identical rules; one for each patrolling profile. When you have created the first rule, you can make a copy of it, and quickly create the second rule based on the copy. Both rules are covered in this example.

Prerequisites

This rule is based on a PTZ camera being able to patrol according to two different patrolling profiles, and a time profile being used to determine which patrolling profile should be used. Before creating a rule like this, always verify the following:

- You have specified a time profile covering at least one of the time periods you want to differentiate between. You could specify time profiles covering both time periods, but it will not be necessary since rules can be set up to apply *within* as well as outside a time profile.

  How to specify a time profile...

  To specify a time profile, expand *Rules and Events* in the Management Client's Site Navigation pane (see "Panes Overview" on page 33), then select ***Time Profiles***. The *Time Profiles* list will appear. In the *Time Profiles* list, right-click ***Time Profiles***, and select ***Add Time Profile...*** For detailed information about specifying time profiles, refer to Manage time profiles (on page 173).

- The camera in question is a PTZ camera.
- Preset positions and at least two patrolling profiles are defined for the camera.

  How to define preset positions and patrolling profiles...

  When a PTZ camera patrols according to a patrolling profile, it moves between a number of preset positions. Thus, before you able to define patrolling profiles for a PTZ camera, the preset positions required for the patrolling profiles must be defined.

  To define preset positions for a PTZ camera, expand *Devices* in the Management Client's Site Navigation pane (see "Panes Overview" on page 33)and select *Cameras*. This will display a list of cameras in the Overview pane (see "Panes Overview" on page 33). Select the required PTZ camera from the list, and select the ***Presets*** tab in the Properties pane (see "Panes Overview" on page 33). For details of how to define preset positions on the *Presets* tab, refer to Preset Positions (see "PTZ Presets tab (camera properties)" on page 96).

  Once you have defined the required preset positions, patrolling profiles for the PTZ camera are defined on the neighboring *Patrolling* tab. For details of how to define patrolling profiles on the *Patrolling* tab, refer to Patrolling (see "PTZ Patrolling tab (camera properties)" on page 93).

Creating the First Rule; Patrolling During Daytime

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Rules and Events*, then right-click Rules and select *Add New Rule...*:

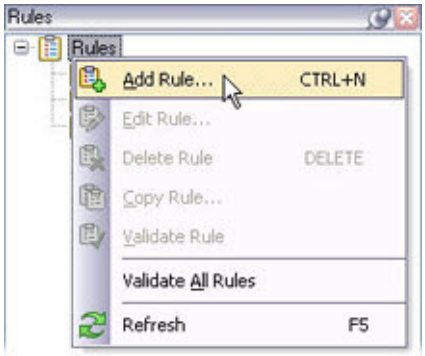

2. The *Manage Rule* wizard opens. Type a name for the new rule in the *Rule name* field.

   In this example...the rule will cover a specific camera and how it should patrol during daytime. We therefore overwrite the default rule name (e.g. *New Rule 001*) with a descriptive name:

   

   Tip: Always use a descriptive name for the rule. Once you have several rules, you will find that descriptive names are a great help when identifying individual rules.

3. On Step 1 of *Manage Rule*, select the required rule type.

   In this example...we want to base the rule on a time period. Therefore, we select *Perform an action in a time interval*:

   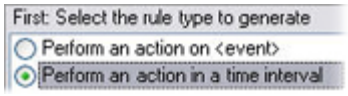

   Click *Next* to go to step 2 of the wizard.

4. On step 2 of the wizard, specify which time conditions should be met in order for the rule to apply.

   In this example... we want the rule to apply within a specific time profile, so we select the time condition *Within selected time in <time profile>*:

   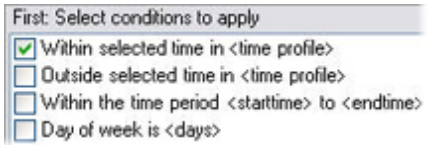

   Based on our selection, the wizard prompts us to specify the required time profile in the rule description:

   

   Click the underlined item to specify the exact content of the rule description.

   In this example... we click the time *profile* link, select the time profile *Daytime*, and click *OK*:

   

   The rule description now reflects our selection:

   

   Click *Next* to move to step 3 of the wizard.

5. On step 3 of the wizard, first specify which actions to perform.

   In this example...we want patrolling according to a specific patrolling profile. We therefore select the action *Start patrolling on <device> using <profile> with PTZ priority <priority>*:

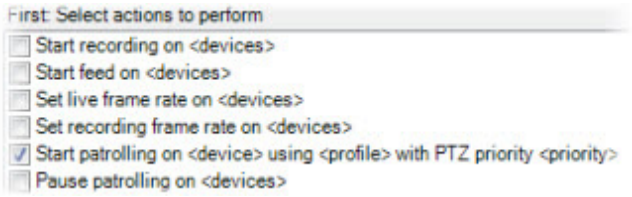

Based on the selection of actions, the wizard extends the rule description, and prompts us to specify the required device, patrolling profile and its priority (see "Actions and Stop actions" on page 136):

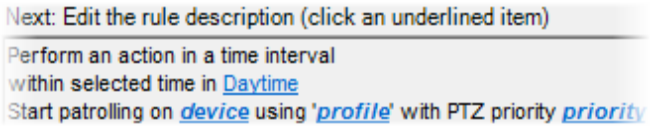

Click the underlined items in the extension of the rule description in order to specify their exact contents:

In this example... we first click the *device* link and in the *Select device* dialog opening we select a device and click *OK*:

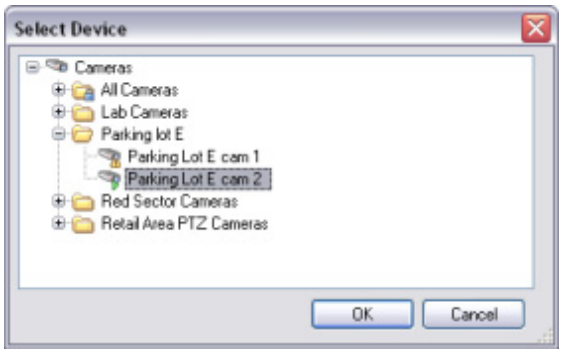

Then we click the *profile* link and select a patrolling profile in the dialog opening and then click OK.

Finally, click the *priority* link to set the priority (see "Actions and Stop actions" on page 136) of the patrolling profile.

The rule description now reflects our selection:

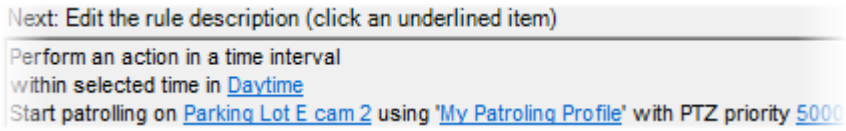

Click *Next* to move to step 4 of the wizard.

6. On step 4 of the wizard, select stop criteria.

   Stop criteria are important in many types of rules. Without a stop criterion, many actions could go on indefinitely once started.

   In this example...Without a stop criterion, the rule in this example would make the PTZ camera start patrolling according to the selected patrolling profile, but it would never stop. Based on the elements in our rule description, we therefore must select a stop criterion. Since our rule is triggered when a time period starts, the wizard automatically suggests that stop action is performed when the time period ends:

The suggestion is also reflected in the rule description. However, we still need to specify exactly which stop action we want performed.

Click *Next* to move to the next step of the wizard.

7. In this step, the wizard suggests one or more stop actions based on the previously selected start actions.

   In this example...Based on the start action *start patrolling* in our rule description, the wizard automatically suggests the stop action *stop patrolling*. It furthermore suggests that patrolling is stopped immediately when the time period ends:

   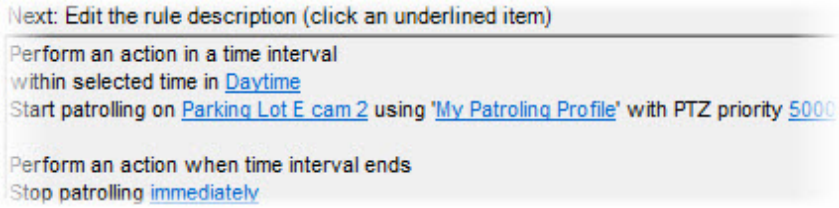

   This is exactly what we want; we do not need to change it.

   Your rule is by default active, meaning that once you have created it, it will be applied as soon as the rule's conditions are met.

   If you do not want the rule to be active straight away, clear the *Active* check box in the top part of the *Manage Rule* window:

   

   Tip: You can always activate/deactivate the rule later.

8. Click *Finish*. This will add your new rule to the list of rules:

   

Creating the Second Rule; Patrolling During Nighttime

Tip: You do not have to create the second rule from scratch; you can make a copy of the first rule, then change it. To copy a rule, select the rule in the *Rules* list, right-click, and select *Copy Rule*... This will open *Manage Rule*, which will display an editable copy of the rule.

1. Copy the first rule, then make the following changes to the rule:
   - Change the rule name so it better describes the new rule, for example to *PTZ Camera Nighttime Patrolling*.
   - On the time conditions selection step, select that the rule should apply not within but *outside* the time profile:

     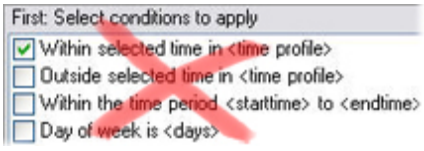

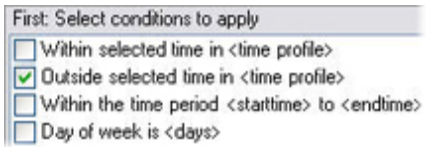

- In the rule description, click the link in the sentence *Start patrolling on ...*, and select a patrolling profile matching your nighttime requirements rather than your daytime requirements:

2. Click *Finish*.

Pause PTZ patrolling and go to PTZ preset on input rule

In this example, we assume that patrolling has already been set up for the PTZ camera, and that the external input unit is a door sensor connected to an input port on a device on the system: When the door sensor is activated, the PTZ camera will pause patrolling, move to a preset position covering the door area, remain at the preset position for 15 seconds, then resume patrolling.

Prerequisites

This rule is based on an input being activated, and on a patrolling PTZ camera moving to a specific preset position. Therefore, an external input unit must be available, i.e. connected to the input port of a device on the system. Furthermore, the preset position to which the PTZ camera should move when the rule is applied must have been defined. Before creating a rule like this, always verify the following:

- An external input unit is successfully connected to an input port on a device, and the states of the input unit (activated/deactivated) work as required.
- The camera in question is a PTZ camera with the required preset positions and patrolling defined.

How to define preset positions and patrolling profiles...

When a PTZ camera patrols according to a patrolling profile, it moves between a number of preset positions. Thus, before you able to define patrolling profiles for a PTZ camera, the preset positions required for the patrolling profiles must be defined.

To define preset positions for a PTZ camera, expand *Devices* in the Management Client's Site Navigation pane (see "Panes Overview" on page 33)and select *Cameras*. This will display a list of cameras in the Overview pane (see "Panes Overview" on page 33). Select the required PTZ camera from the list, and select the ***Presets*** tab in the Properties pane (see "Panes Overview" on page 33). For details of how to define preset positions on the *Presets* tab, refer to Preset Positions (see "PTZ Presets tab (camera properties)" on page 96).

Once you have defined the required preset positions, patrolling profiles for the PTZ camera are defined on the neighboring *Patrolling* tab. For details of how to define patrolling profiles on the *Patrolling* tab, refer to Patrolling (see "PTZ Patrolling tab (camera properties)" on page 93).

Creating the Rule

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Rules and Events*, then right-click Rules and select *Add New Rule...*:

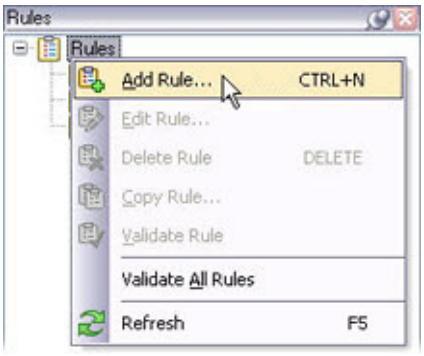

2. The *Manage Rule* wizard opens. Type a name for the new rule in the *Rule name* field.

   In this example... the rule will cover a specific camera (*simply called PTZ Camera*) and how it should behave upon an activated input. We therefore overwrite the default rule name (e.g. *New Rule 001*) with a descriptive name:

   

   Tip: Always use a descriptive name for the rule. Once you have several rules, you will find that descriptive names are a great help when identifying individual rules.

3. On Step 1 of *Manage Rule*, select the required rule type.

   In this example... we want to base the rule on an event. Therefore, we select *Perform an action* on <*event*>. Our selection is immediately reflected in the initial rule description in the lower half or the wizard window:

   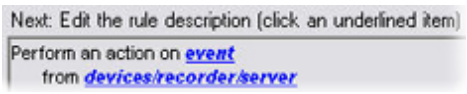

   Click the underlined items in the rule description in order to specify its exact content:

   Event link: Clicking the *event* link lets you select the event which must occur in order for the rule to apply. In order for you to get a good overview, selectable events are listed in groups according to whether they are related to plug-ins, dependent on hardware configuration, built into the system itself, etc.

   In this example...we want the event to be activated input. Input comes from— and is configured on— separate hardware rather than on the system itself, so we go to the *Custom Events* group, select the event *Input Activated*, and click *OK.*

   Devices/recording server/management server link: When you have selected the required event, clicking the *devices/recording server/management server* link opens the *Select Devices and Groups* window, which lets you specify the devices on which the event should occur in order for the rule to apply.

   In this example... the event should occur on an input called *Back Door Sensor* in order for the rule to apply. In the *Select Devices and Groups* window we therefore drag the input *Back Door Sensor* to the *Selected* list and click *OK*. By doing this we have specified the exact content of the first part of the wizard's rule description, which now looks like this:

   

   Click *Next* to move to step 2 of the wizard.

4. On step 2 of the wizard, specify which time conditions should be met in order for the rule to apply.

In this example...we want the rule to apply whenever input is activated on the back door sensor, regardless of time. When creating event-based rules it is possible to bypass the time conditions; we therefore want to skip the wizard's step 2 entirely.

Click *Next* to move to step 3 of the wizard.

5. On step 3 of the wizard, first specify which actions to perform.

In this example...we want to two things to happen: patrolling should pause, and the PTZ camera should move to a specific preset position with a specific priority (see "Actions and Stop actions" on page 136). We therefore select the actions Pause patrolling on *<devices>* and *Move <device> to <preset> position with PTZ priority <priority>*.

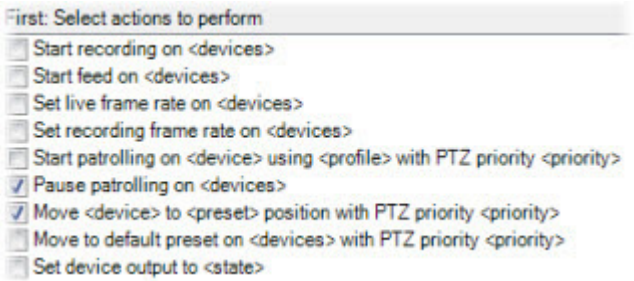

Based on the selection of actions, the wizard automatically extends the rule description in the lower part of the wizard window.

In this example...Based on our selections Pause patrolling on *<devices>* and *Move <device> to <preset> position with PTZ priority <priority>* the wizard automatically suggests an extension to the existing rule description:

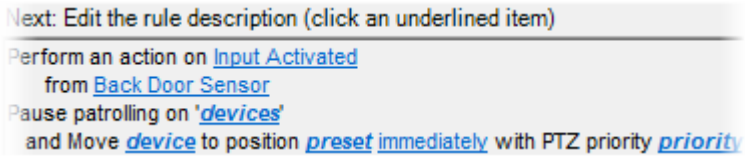

6. Click the underlined items in the extension of the rule description in order to specify its exact content:

devices: Clicking the devices link lets you select the devices on which patrolling should be paused. Only PTZ cameras will be selectable.

In this example...patrolling should be paused on our PTZ camera. In the *Select Group Members* window we therefore drag *PTZ Camera to the Selected* list and click *OK*.

device: Clicking the *device* link lets you select to move another device than the device(s) on which patrolling was paused. You are also able to select to move the device on which patrolling was paused.

In this example...we select to move the same PTZ camera on which patrolling was paused, and click OK:

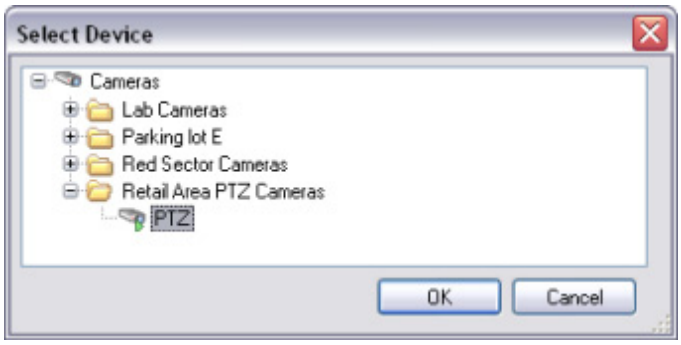

preset: Clicking the *preset* link lets you select which preset position the PTZ camera should move to. You will be able to select from a list of preset positions defined for the PTZ camera you selected before.

In this example...we select a preset position called *Back Door*, and click *OK*:

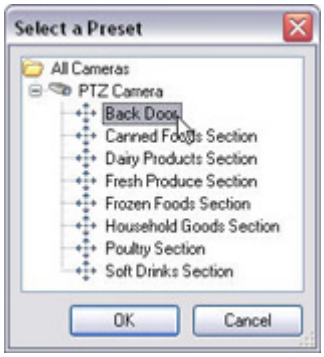

immediately: The wizard automatically suggests that the camera moves to the preset position *immediately* after it has paused patrolling. Clicking the immediately link lets you specify a delay, if required.

priority: Clicking the *priority* link lets you specify the priority (see "Actions and Stop actions" on page 136) of the camera position.

In this example...the wizard's suggestion *immediately* suits us fine, so we leave it as it is.

The rule description now indicates which camera will pause patrolling, which preset position it will move to, and how soon:

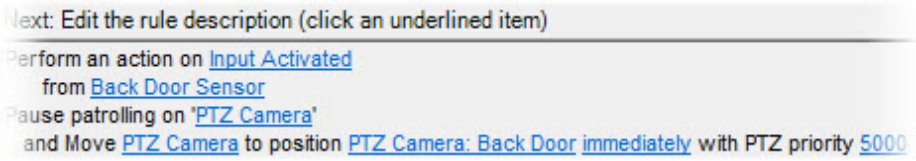

Click *Next* to move to step 4 of the wizard.

7. On step 4 of the wizard, select stop criteria.

   Stop criteria are important in many types of rules. Without a stop criterion, many actions could go on indefinitely once started.

   In this example...Without a stop criterion, the rule in this example would make the camera pause patrolling, then move to the preset position and remain there indefinitely. Based on the elements in our rule description, we therefore **must** select a stop criterion.

   Since our rule is triggered by an event, the wizard automatically suggests that we base our stop action on an event as well. In the rule description, the wizard even suggests that the stop action is performed when input is

deactivated on the back door sensor. However, we want something different, so we select *Perform stop action after <time>*:

In the rule description, the wizard now prompts us to specify the required time:

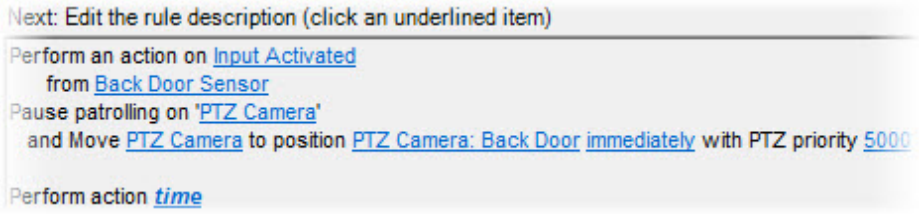

We click the *time* link, specify 15 seconds, and click *OK*:

The rule description now indicates the 15 seconds selected.

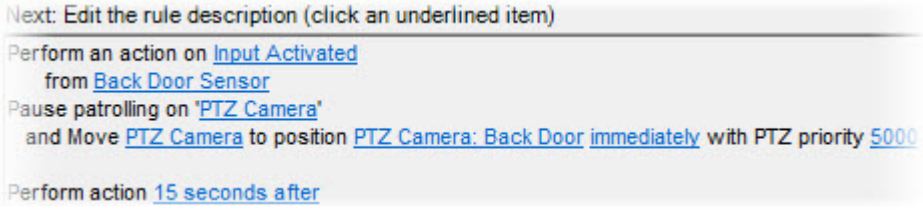

Based on the start action *pause patrolling* in our rule description, the wizard automatically suggests the stop action *resume patrolling*. It furthermore suggests that patrolling is resumed immediately after the 15 second pause:

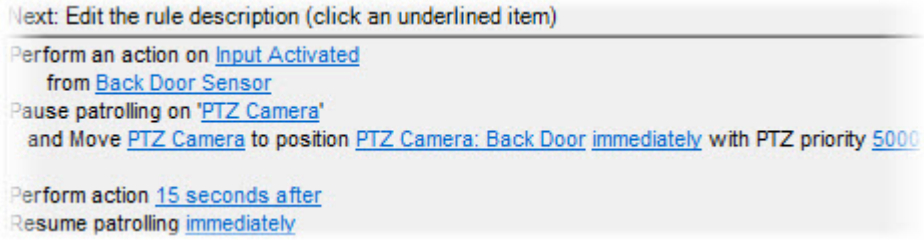

This is exactly what we want; we do not need to change it, although by clicking the *immediately* link we could have specified a delay of e.g. 3 seconds.

Your rule is by default active, meaning that once you have created it, it will be applied as soon as the rule's conditions are met. If you do not want the rule to be active straight away, clear the *Active* check box in the top part of the *Manage Rule* window:

Tip: You can always activate/deactivate the rule later.

8. Click *Finish*. This will add your new rule to the list of rules:

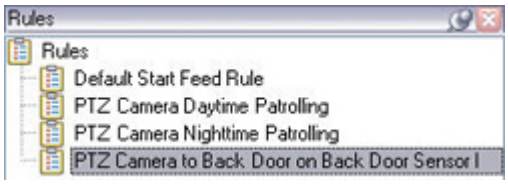

Default rules

Your system comes with a number of default rules, ensuring that basic features work without any user intervention being required.

IMPORTANT: Like other rules, default rules can be deactivated and/or modified as required. The fact that default rules are present does therefore not in itself guarantee that your system will work as desired. Nor does it guarantee that video feeds or audio feeds will automatically be fed to the system, as the default rules may subsequently have been deactivated or modified.

Default goto preset when PTZ is done rule

Ensures that PTZ (Pan/Tilt/zoom) cameras go to their respective default preset positions after they have been operated manually.

IMPORTANT: This rule is by default not enabled. Even when the rule is enabled, you must have defined default preset positions for the required PTZ cameras in order for the rule to work; you do this on the *Presets* tab (see "PTZ Presets tab (camera properties)" on page 96).

In case you accidentally delete the default goto preset when PTZ is done rule, you can recreate it with the following content:

```
Perform an action on PTZ Manual Session Stopped from All Cameras

Move immediately to default preset on the device on which event occurred
```

Default record on motion rule

Ensures that as long as motion is detected in video from cameras, the video is recorded, provided recording is enabled (see "Record tab overview" on page 125) for the cameras in question (recording is by default enabled).

IMPORTANT: While the default rule specifies recording based on detected motion, it does not guarantee that video will be recorded, as individual cameras' recording may have been disabled for one or more cameras. Even when recording is enabled, bear in mind that the quality of recordings may be affected by individual camera's recording settings.

In case you accidentally delete the default record on motion rule, you can recreate it with the following content:

```
Perform an action on Motion Started from All Cameras start recording 3 seconds
before on the device on which event occurred

Perform stop action on Motion Stopped from All Cameras stop recording 3 seconds
after
```

Default record on request rule

Ensures that video is recorded automatically when an external request occurs, provided recording is enabled (see "Record tab overview" on page 125) for the cameras in question (recording is by default enabled).

IMPORTANT: The request is always triggered by a system integrating externally with your system, and the rule is primarily used by integrators of external systems or plug-ins. In case you accidentally delete the default record on bookmark rule, you can recreate it with the following content:

```
Perform an action on Request Start Recording from External start recording immediately on the devices from metadata

Perform stop action on Request Stop Recording from External stop recording immediately
```

Default start audio feed rule

Ensures that audio feeds from all connected microphones and speakers are automatically fed to the system.

IMPORTANT: While the default rule enables access to connected microphones' and speakers' audio feeds immediately upon installing the system, it does not guarantee that audio will be recorded (see "Record tab overview" on page 125), as recording settings must be specified separately.

In case you accidentally delete the default start audio feed rule, you can recreate it with the following content:

```
Perform an action in a time interval always start feed on All Microphones, All Speakers

Perform an action when time interval ends stop feed immediately
```

Default start feed rule

Ensures that video feeds from all connected cameras are automatically fed to the system.

IMPORTANT: While the default rule enables access to connected cameras' video feeds immediately upon installing the system, it does not guarantee that video will be recorded, as cameras' recording settings must be specified separately.

In case you accidentally delete the default start feed rule, you can recreate it with the following content:

```
Perform an action in a time interval always start feed on All Cameras

Perform an action when time interval ends stop feed immediately
```

Events overview

Events may be created on the recording server or in Ocularis Base. The information discussed here refer to setting events on the recorder.

When creating an event-based rule in the *Manage Rule* wizard (see "Manage rules" on page 165), you are able to select between a number of different events.

In order for you to get a good overview, selectable events are listed in groups according to whether they are:

Some hardware is capable of creating events themselves, for example to detect motion. These can be used as events but must obviously be configured on the hardware before they can be used in the system. Events listed here may only be possible on some hardware. For example, only selected cameras are able to detect tampering or temperature changes.

Configurable events, hardware

These configurable events are unknown until they are automatically imported from device drivers. As a result, they cannot be documented separately or in details in this context. Furthermore, configurable events are not triggered until they have been added and configured on the ***Event*** tab on a hardware or device (see "Events tab overview" on page 128).

Predefined events, hardware

Name	Description
Communication Error (Hardware)	Occurs when a connection to a the hardware is lost.
Communication Started (Hardware)	Occurs when communication with the hardware is successfully established.
Communication Stopped (Hardware)	Occurs when communication with the hardware is successfully stopped.

Configurable events, devices

These configurable events are unknown until they are automatically imported from device drivers. As a result, they cannot be documented separately or in details in this context. Furthermore, configurable events are not triggered until they have been added and configured on the ***Event*** tab on a hardware or device (see "Events tab overview" on page 128).

Predefined events, devices

Name	Description
Bookmark Reference Requested	N/A
Communication Error (Device)	Occurs when a connection to a device is lost; or when an attempt is made to communicate with a device, and the attempt is unsuccessful.
Communication Started (Device)	Occurs when communication with a device is successfully established.
Communication Stopped (Device)	Occurs when communication with a device is successfully stopped.
Feed Overflow Started	Feed overflow (a.k.a. Media overflow) occurs when a recording server is unable to process received video as quickly as specified in the configuration and therefore is forced to discard some images. If the server is healthy, feed overflow usually happens because of slow disk writes. It can be resolved either by reducing the amount of data written, or by improving the storage system's performance. Reduce the amount of written data by reducing frame rates, resolution or image quality on your cameras. This will in general degrade recording quality. If you are not interested in that, instead improve your storage system's performance by installing extra drives to share the load or by installing faster disks or controllers. **Tip:** This rare event can be used for triggering actions that will help you avoid the problem, e.g. for lowering the recording frame rate.
Feed Overflow Stopped	Occurs when feed overflow (see description of the *Feed Overflow Started* event) ends.
Live Client Feed Requested	Occurs when client users request a live stream from a device. The event occurs upon the request— even if the client user's request subsequently turns out to be unsuccessful, for example because the client user does not have the rights required for viewing the requested live feed or because the feed is for some reason stopped.
Live Client Feed Terminated	Occurs when client users no longer request a live stream from a device.

Name	Description
Motion Started	Occurs when the system detects motion in video received from cameras. This type of event requires that the system's motion detection is enabled for the cameras to which the event will be linked. Exactly what constitutes motion depends on the motion detection settings specified for individual cameras in the system. In addition to the system's motion detection, some cameras are— depending on configuration of the camera hardware— themselves able to detect motion. Such camera-detected motion detection can also be used in system rules, however they do not work until configured on the camera hardware itself. Refer to Configurable events, devices (on page 162).
Motion Stopped	Occurs when motion is no longer detected in received video. See also the description of the *Motion Started* event.
Output Activated	Occurs when an external output unit connected to an output port on a device is activated. This type of event requires that at least one device on your system has an external input unit connected to an output port.
Output Changed	Occurs when the state of an external output unit connected to an output port on a device is changed, regardless of which state the external input unit is changed to. This type of event requires that at least one device on your system has an external input unit connected to an output port.
Output Deactivated	Occurs when an external output unit connected to an output port on a device is deactivated. This type of event requires that at least one device on your system has an external input unit connected to an output port.
PTZ Manual Session Started	Occurs when a manually operated PTZ session (as opposed to a PTZ session based on scheduled patrolling or automatically triggered by an event) is started on a camera. This type of event requires that the cameras to which the event will be linked are PTZ (Pan/Tilt/Zoom) cameras.
Manual PTZ Session Stopped	Occurs when a manually operated PTZ session (as opposed to a PTZ session based on scheduled patrolling or automatically triggered by an event) is stopped on a camera. This type of event requires that the cameras to which the event will be linked are PTZ (Pan/Tilt/Zoom) cameras.
Motion Stopped	Occurs when motion is no longer detected in received video. See also the description of the *Motion Started* event.
Recording Started	Occurs when recording is started.
Recording Stopped	Occurs when recording is stopped.
Settings Changed	Occurs when settings on a device are successfully changed.
Settings Changed Error	Occurs when an attempt is made to change settings on a device, and the attempt is unsuccessful.

Predefined events, external

Name	Description
Request Start Recording	N/A
Request Stop Recording	N/A

User-defined events, external

A number of events custom made to suit your system may also be selectable. Such user-defined events can be used for:

- Making it possible for end users to manually trigger events while viewing live video in the Ocularis Client.
- Countless other purposes. For example, you may create user-defined events which will occur if a particular type of data is received from a device.

For information about how to define user-defined events in the Management Client, refer to Manage user-defined events (on page 180).

Recording servers

Name	Description
Archive Available	Occurs when an archive (see "About storage and archiving" on page 61) for a recording server becomes available after having been unavailable (see *Archive Unavailable* next).
Archive Unavailable	Occurs when an archive (see "About storage and archiving" on page 61) for a recording server becomes unavailable, for example if the connection to an archive located on a network drive is lost. When this is the case, it will not be possible to archive recordings. You can use the event to, for example, trigger a notification profile so an e-mail notification is automatically sent to relevant people in your organization.
Archive Not Finished	Occurs when an archive (see "About storage and archiving" on page 61) for a recording server is not finished with the last archiving round when the next is scheduled to start.
Database Disk Full	Occurs when a database disk is full. A database disk is considered to be full when there is less than 5GB of space is left on the disk: The oldest data in a database will always be auto-archived (or deleted if no next archive is defined) when less than 5GB of space is free. If less than 1GB space is free, data will be deleted. A database always requires 250MB of free space; if this limit is reached (if data is not deleted fast enough), no more data will be written to the database until enough space has been freed. The actual maximum size of your database will thus be the amount of gigabytes you specify, minus 5GB.
Database Full - Auto Archive	Occurs when an archive (see "About storage and archiving" on page 61) for a recording server is full and needs to auto-archive to an archive in the hierarchy.
Database Repair	Occurs if a database becomes corrupted, in which case the system will automatically attempt two different database repair methods: a fast repair and a thorough repair.

Name	Description
Database Storage Area Available	Occurs when a storage area (see "About storage and archiving" on page 61) for a recording server becomes available after having been unavailable (see *Database Storage Area Unavailable* next). You can, for example, use the event to start recording if it has been stopped by a *Database Storage Area Unavailable* event.
Database Storage Area Unavailable	Occurs when a storage area (see "About storage and archiving" on page 61) for a recording server becomes unavailable, for example if the connection to a storage area located on a network drive is lost. When this is the case, it will not be possible to store recordings. You can use the event to, for example, stop recording and trigger a notification profile (see "Manage notification profiles" on page 176) so an e-mail notification is automatically sent to relevant people in your organization.
Failover Started	Occurs when a failover recording server (see "About failover recording servers—regular and hot standby" on page 234) takes over from a recording server. A failover recording server is a spare recording server which can take over if a standard recording server becomes unavailable.
Failover Stopped	Occurs when a recording server becomes available again, and is able to take over from a failover recording server (see "About failover recording servers—regular and hot standby" on page 234).

Manage rules

Rules are a central element in your system. Rules determine highly important settings, such as when cameras should record, when PTZ (Pan/Tilt/Zoom) cameras should patrol, when notifications should be sent, etc.

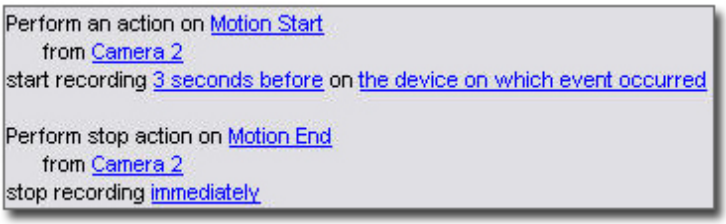

Example: A rule specifying that a particular camera should begin recording when it detects motion

You create and manage rules in the Management Client.

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand the *Rules and Events* folder, then select *Rules*. In the Overview pane (see "Panes Overview" on page 33), a *Rules* list, providing an overview of all existing rules in your system, will appear.
2. If needed, create, edit, copy and/or validate rules from the list.

Tip: You do not necessarily have to delete an unwanted rule; you may also just temporarily deactivate the rule by clearing the *Active* check box in the *Rule Information* pane for the rule in question, then saving the setting by selecting the Management Client's *File* menu.

Show me where to find the ***Active*** check box

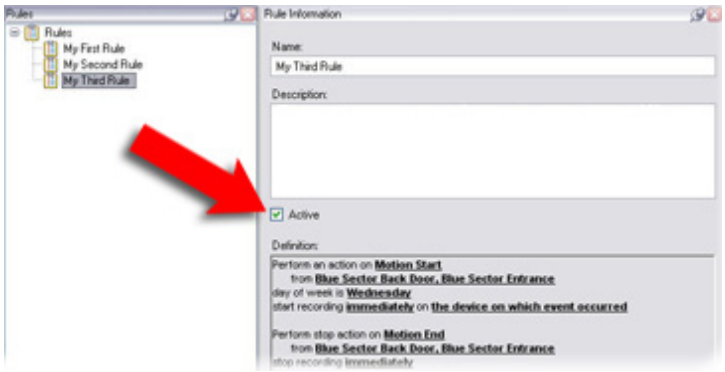

About rules (on page 166)

About rules

In short, rules specify actions which should be carried out under particular conditions. Example: When motion is detected (condition), a camera should begin recording (action).

The following are *examples* of what you can do with rules:

- Start and stop recording
- Set non-default live frame rate
- Set non-default recording frame rate
- Start and stop PTZ patrolling
- Pause and resume PTZ patrolling
- Move PTZ cameras to specific positions
- Set output to activated/deactivated state
- Send notifications via e-mail
- Generate log entries
- Generate events
- Apply new device settings, for example a different resolution on a camera
- Start and stop plug-ins
- Start and stop feeds from devices

How is stopping the feed from a device different from manually disabling the device? Stopping a device means that video will no longer be transferred from the device to the system, in which case neither live viewing nor recording will be possible. However, a device on which the feed has been stopped will still be able to communicate with the recording server, and the feed from device can be started automatically through a rule, as opposed to when the device is manually disabled in the Management Client.

IMPORTANT: Some rule content may require that certain features are enabled for the devices in question. For example, a rule specifying that a camera should record will not work as intended if recording is not enabled for the camera in question. Before creating a rule it is therefore highly recommended that you verify that the devices involved will be able to perform as intended. For a number of typically required rules, such prerequisites are described in Create typical rules (on page 142).

How a rule is triggered

Two types of conditions can trigger rules:

Name	Description
Events	When events occur on the surveillance system (for example when motion is detected, when the system receives input from external sensors, etc.)
Time	When specific periods of time are entered (for example *Thursday 16th August 2007 from 07.00 to 07.59*, or *every Saturday and Sunday*.)

What you can cover in a rule

Your exact number of options depends on the type of rule you want to create, and on the number of devices available on your system.

Rules, however, provide a high degree of flexibility: You are able to combine event and time conditions, you are able to specify several actions in a single rule, and very often you are able to create rules covering several or all of the devices on your system.

You can make your rules as simple or complex as required. For example, you can create very simple time-based rules:

> **Example** **Very Simple Time-Based Rule:** On Mondays between 08.30 and 11.30 (time condition), Camera 1 and Camera 2 should start recording (action) when the time period begins and stop recording (stop action) when the time period ends.

And you can create very simple event-based rules, involving events on one device only:

> **Example** **Very Simple Event-Based Rule:** When motion is detected (event condition) on Camera 1, Camera 1 should start recording (action) immediately, then stop recording (stop action) after 10 seconds.

However, even though an event-based rule is activated by an event on one device, you can specify that actions should take place on one or more other devices.

> **Example** **Rule Involving Several Devices:** When motion is detected (event condition) on Camera 1, Camera 2 should start recording (action) immediately, and the siren connected to Output 3 should sound (action) immediately; then, after 60 seconds, Camera 2 should stop recording (stop action), and the siren connected to Output 3 should stop sounding (stop action).

You can of course also combine events and scheduled times in a rule:

> **Example** **Rule Combining Time, Events, and Devices:** When motion is detected (event condition) on Camera 1, and the day of the week is Saturday or Sunday (time condition), Camera 1 and Camera 2 should start recording (action) immediately, and a notification should be sent to the security manager (action); then, 5 seconds after motion is no longer detected on Camera 1 or Camera 2, the 2 cameras should stop recording (stop action).

The required complexity of rules will vary from organization to organization: Some may require only a number of simple rules; some may require a mix of simple and complex rules.

Create many simple or a few complex rules?

Depending on your organization's requirements, it is often a good idea to create many simple rules rather than a few complex rules. Even though this will lead to you having more rules, it generally makes it much easier for you to maintain an overview of what your rules do.

Keeping your rules simple also means that you have much more flexibility when it comes to deactivating/activating individual rule elements— with simple rules, you can deactivate/activate entire rules when required.

Default rules

Your system comes with a number of default rules (on page 160), ensuring that basic features work without any user intervention being required.

Add a rule

When you create rules, you are guided by the wizard *Manage Rule* which provides a highly intuitive approach. It helps you stay focused by listing only relevant options. It ensures that a rule will not contain missing elements. And, finally, based on your rule's content, it automatically suggests suitable stop actions (i.e. what should take place when the rule no longer applies), ensuring that you do not unintentionally create a never-ending rule.

1. In the Overview pane (see "Panes Overview" on page 33), right-click the *Rules* item, and select *Add Rule...*:

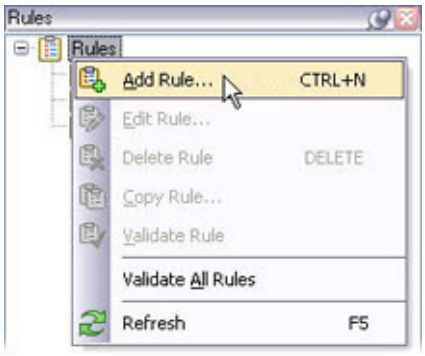

This will open the wizard *Manage Rule*:

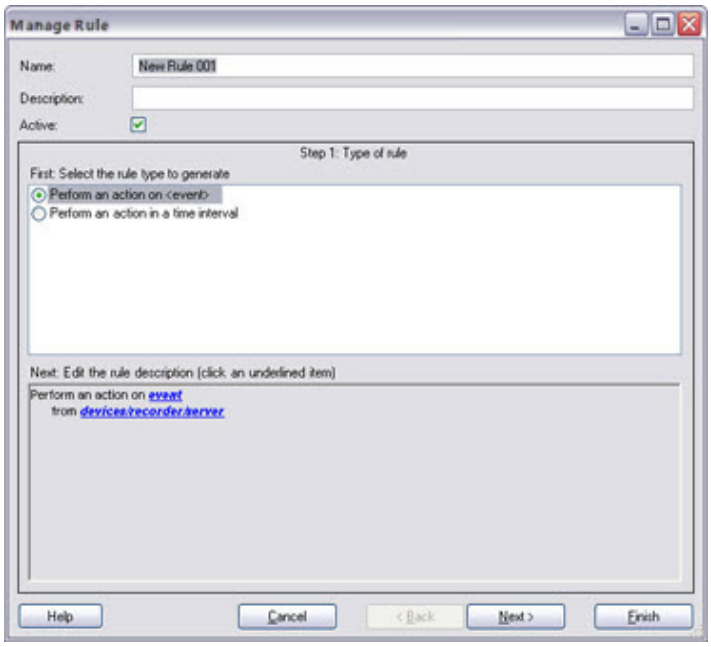

The wizard guides you through the process of specifying the content of your rule. The wizard makes the process interactive, yet intuitive: based on your main selections, it asks you to specify your exact requirements for the rule.

2. Begin by specifying a name (compulsory) and a description (optional) of the new rule in the *Name* and *Description* fields respectively.

   Tip: Always use a descriptive name for the rule. Once you have several rules, you will find that descriptive names are a great help when identifying individual rules.

3. Then select the required type of condition for the rule: either a rule which performs one or more actions when a particular event occurs, or a rule which performs one or more actions when a specific period of time is entered:

   Perform an action on <event>

   If you select an event-based rule, the lower part of the wizard window will display an initial rule description:

   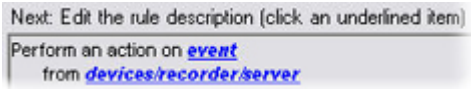

   Click the underlined items in the rule description in order to specify its exact content:

   - ***Event***: Clicking the *event* link lets you select the event which must occur in order for the rule to apply (for example *Motion Started*).
   - ***Devices/recording server/management server:*** When you have selected the required event, clicking the *devices/recording server/management server* link lets you specify the devices on which the event should occur in order for the rule to apply. Depending on your event specification, you may be able to select from a list of cameras, inputs, outputs, etc. In this example illustration, the selectable devices are all cameras:

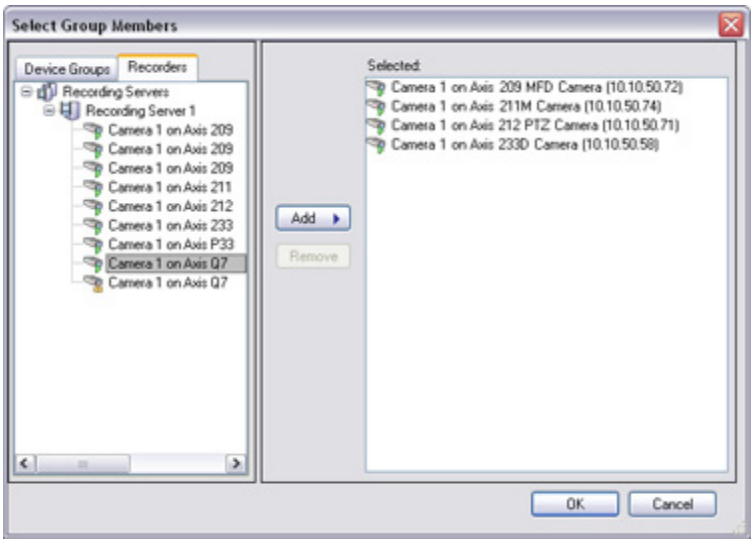

You specify the required devices by moving them from the *Available devices* list to the *Selected devices* list.

To move a device from the *Available devices* list to the *Selected devices* list, either select the device and click the *Add* button, double-click the device, or simply drag the device from one list to the other.

Tip: When devices are grouped into so-called device groups, you can quickly move all devices in a group simply by moving the group folder.

When the required devices are listed in the *Selected devices* list, click *OK*.

You have now specified the exact content of the first part of the rule description:

Example only; your selections may be different

Perform an action in a time interval

If you select a time-based rule, no more information is required on the wizard's first step.

4. Click *Next* to go to the wizard's second step. On the wizard's second step you are able to define further conditions for the rule.
5. Select one or more conditions, for example *Day of week is <day>*:

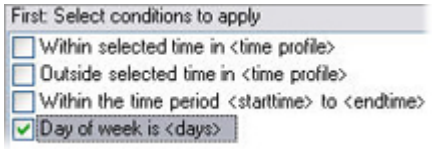

Example only; your selections may be different

Depending on your selections, the lower part of the wizard window lets you edit the rule description:

Example only; your selections may be different

Click the underlined items in ***bold italics*** to specify their exact content. For example, clicking the *days* link in our example would let you select one or more days of the week on which the rule should apply.

6. Having specified your exact conditions, click *Next* to move to the next step of the wizard and select which actions should be covered by the rule.

   Depending on the content and complexity of your rule, further wizard steps may let you define further information, such as stop events and stop actions. For example, if a rule specifies that a device should perform a particular action during a time interval (for example Thursday between 08.00 and 10.30), the wizard may ask you to specify what should happen when that time interval ends.

7. Your rule is by default active, meaning that once you have created it, it will be applied as soon as the rule's conditions are met.

   If you do not want the rule to be active straight away, clear the *Active* check box:

   Tip: You can always activate/deactivate the rule later.

8. Click *Finish*.

To view step-by-step descriptions of how to create typically required rules, refer to Create typical rules (on page 142).

Edit, copy and rename a rule

1. In the Overview pane (see "Panes Overview" on page 33), right-click the required rule.
2. Select either:

   Edit Rule...

   or

   Copy Rule...

   or

   Rename Rule...

   depending on your needs.

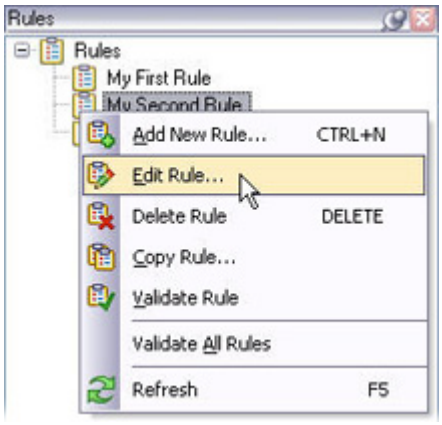

   Example when selecting *Edit Rule...*

   The wizard *Manage Rule* opens.

3. In the wizard, rename and/or change the rule as required. If you selected *Copy Rule...*, the wizard opens, displaying a copy of the selected rule.

4. Click *Finish*.

Deactivate and activate a rule

A rule is by default active, meaning that your system applies the rule as soon as the rule's conditions apply. If you do not want a rule to be active, you can deactivate the rule. When the rule is deactivated, the system does not apply the rule, even if the rule's conditions apply. A deactivated rule can easily be activated later.

Deactivating a rule:

1. In the Overview pane (see "Panes Overview" on page 33), select the required rule.
2. Clear the *Active* check box in the Properties pane (see "Panes Overview" on page 33):

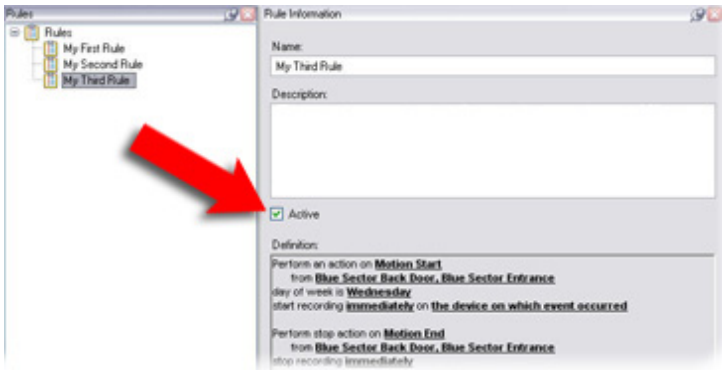

3. Save the setting by clicking *Save* in the Management Client's toolbar (see "Management Client Overview" on page 30).
4. The deactivated rule will be indicated by a different icon in the *Rules* list:

Example: Different icon indicates that third rule is deactivated

Activating a rule:

When you want to activate the rule again, select the required rule, select the *Activate* check box, and save the setting.

Validate rule(s)

You are able to validate the content of an individual rule or all rules in one go.

Why would I need to validate the content of rules? When you create a rule, the *Manage Rule* ensures that all of the rule's elements make sense. However, when a rule has existed for some time, one or more of the rule's elements may have been affected by other configuration, and the rule may no longer work. For example, if a rule is triggered by a particular time profile, the rule will not work if the time profile in question has subsequently been deleted. Such unintended effects of configuration may be hard to keep an overview of; rule validation helps you keep track of which rules have been affected.

IMPORTANT: Validation takes place on a per-rule basis; each rule is validated in isolation. It is currently not possible to validate rules against each other (for example in order to see whether one rule conflicts with another rule), not even if using the *Validate All Rules* feature.

Furthermore, it is not possible to validate whether configuration of prerequisites outside the rule itself may prevent the rule from working. For example, a rule specifying that recording should take place when motion is detected by a

particular camera will validate OK if the elements in the rule itself are correct, even though motion detection (which is enabled on a camera level, not through rules) has not been enabled for the camera in question.

To validate an individual rule or all rules in one go, do the following in the Management Client:

1. In the Overview pane (see "Panes Overview" on page 33), right-click the rule you wish to validate, and select *Validate Rule* or *Validate All Rules* (depending on your needs):
2. A simple dialog will inform you whether the rule(s) validated successfully or not. If you chose to validated more than one rule and one or more rules did not succeed, the dialog will list the names of the affected rules:

Manage time profiles

Time profiles are periods of time defined by the administrator. Time profiles can be used when creating rules (see "Manage rules" on page 165), for example, a rule specifying that a certain action should take place within a certain time period. As an alternative to time profiles, refer to Day length time profiles (see "Manage day length time profiles" on page 176).

Time profiles are also assigned to roles (see "Manage roles" on page 184). Per default, all roles are assigned the default time profile *Always.* This means that members of roles with this default time profile attached has no time-based limits to their user rights in the system. An alternative time profile can easily be assigned to a role (see "Add a role" on page 185).

Time profiles are highly flexible: they can be based on one or more single periods of time, on one or more recurring periods of time, or a combination of single and recurring times. Many users will be familiar with the concepts of single and recurring time periods from calendar applications, such as the one in Microsoft® Outlook.

Time profiles always apply in local time. This means that if your system has recording servers placed in different time zones, any actions (e.g. recording on cameras) associated with time profiles will be carried out in each recording server's local time. Example: If you have a time profile covering the period 08.30 to 09.30, any associated actions on a recording server placed in New York will be carried out when the local time is 08.30 to 09.30 in New York, while the same actions on a recording server placed in Los Angeles will be carried out some hours later, when the local time is 08.30 to 09.30 in Los Angeles.

You create and manage time profiles in the Management Client by expanding the Site Navigation pane (see "Panes Overview" on page 33)'s *Rules and Events* folder, then selecting *Time Profiles*. A *Time Profiles* list will appear in the Overview pane (see "Panes Overview" on page 33):

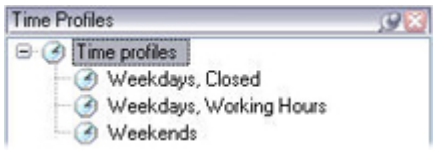

Example only

Specify a time profile

1. In the *Time Profiles* list, right-click *Time Profiles*, and select *Add Time Profile...*. This will open the *Time Profile* window:

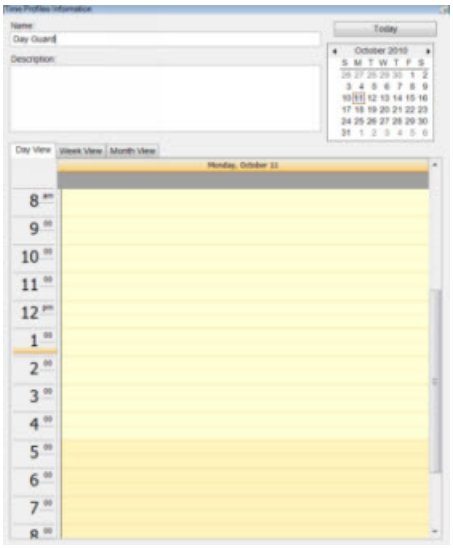

Time and date format may be different on your system

2. In the *Time Profile* window, type a name for the new time profile in the *Name* field. Optionally, type a description of the new time profile in the *Description* field.
3. In the *Time Profile* window's calendar, select either *Day View, Week View* or *Month View*, then right-click inside the calendar and select either *Add Single Time...* or *Add Recurrence Time...*

Tip: If you select a time period by dragging in the calendar before right-clicking, the selected period will automatically be used in the dialog that appears when you select *Add Single Time...* or *Add Recurring Time...*

Specify a single time

When you select *Add Single Time...*, the *Select Time* window appears:

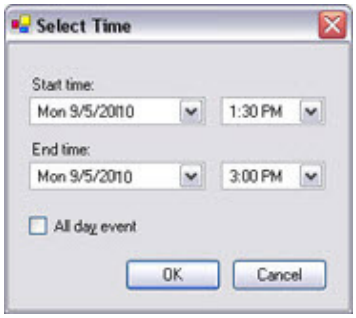

Time and date format may be different on your system

1. In the *Select Time* window, specify *Start time* and *End time*. If the time is to cover whole days, select the *All day event* box.
2. Click *OK*.

Tip: A time profile is able to contain several periods of time. If you want your time profile to contain further periods of time, add more single times or recurring times.

Specify a recurring time

When you select *Add Recurring Time...*, the *Select Recurring Time* window appears:

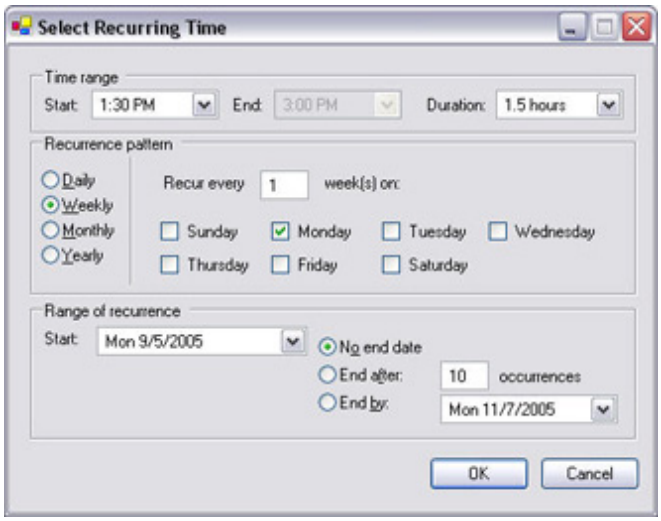

Time and date format may be different on your system

1. In the *Select Time* window, specify time range, recurrence pattern and range of recurrence.
2. Click *OK*.

Tip: A time profile is able to contain several periods of time. If you want your time profile to contain further periods of time, add more single times or recurring times.

1. When you have specified the required time periods for your time profile, click *OK* in the *Time Profile* window. Your new time profile is added to the *Time Profiles* list in the Overview pane (see "Panes Overview" on page 33).

   If at a later stage you wish to edit or delete the time profile, you can do that from the *Time Profiles* list.

Edit a time profile

1. In the Overview pane (see "Panes Overview" on page 33)'s *Time Profiles* list, right-click the required time profile, and select *Edit Time Profile...*.

   Tip: Instead of right-clicking to select *Edit Time Profile*, you can select the required time profile and press F2 on your keyboard.

   This will open the *Time Profile* window.

2. In the *Time Profile* window, edit the time profile as required.

   When you have made the required changes to the time profile, click *OK* in the *Time Profile* window. You will be returned to the Overview pane's *Time Profiles* list.

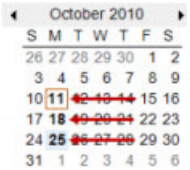

You browse months by clicking the small back/forward buttons.

Tip: In the *Time Profile Information* window, edit the time profile as required. Remember that a time profile may contain more than one time period, and that time periods may be recurring.

Tip: The small month overview in the top right corner of the *Time Profile Information* window can help you get a quick overview of the time periods covered by the time profile, as dates containing specified times are highlighted in bold. In this example, the bold dates indicate that time periods have been specified on several days, and that a recurring time may have been specified on Mondays.

Manage day length time profiles

When cameras are placed outside, it is often required to lower the cameras resolution, enable black/white, or change other settings when it gets dark or vice versa when it gets light. The further north or south from the equator the cameras are placed, the more the sunrise and sunset time varies during the year. This makes it impossible to use normal fixed time profiles (see "Manage time profiles" on page 173) to adjusts camera settings according to light conditions.

To overcome this, day length time profiles can be created and defined in the system according to the sunrise and sunset in a specified geographical area. Via GPS coordinates, the system, on a daily basis, calculates the sunrise and sunset time, even incorporating daylight saving time. As a result, it automatically follows the yearly changes in sunrise/sunset in the selected area, ensuring the profile to be active only when needed. All times and dates are based on the management servers time and date settings.

In addition, you can set a positive or negative offset (in minutes) for the start (sunrise) and end time (sunset). The offset for the start and the end time can be identical or different.

Day length time profiles can be used when creating both rules (see "Manage rules" on page 165) and roles (see "About roles" on page 182).

Create a day length time profile

1. In the Management Client, expanding the Site Navigation pane (see "Panes Overview" on page 33)'s *Rules and Events* folder, select *Time Profiles*.
2. In the Overview pane (see "Panes Overview" on page 33), in the *Time Profiles* list, right-click *Time Profiles*, and select *Add Day Length Time Profile...*.
3. In the *Day Length Time Profile* window, fill in the needed information. In order to deal with transition periods between lightness and darkness, it is possible to offset activation and deactivation of the profile.

   Also, time and month names are shown in the language dictated by your computer's language/regional settings.
4. To see the location of the entered GPS coordinates in a map, click *Show Position in Browser...* (will open a browser).
5. Click *OK*.

Day length time profile properties

Set the following properties for day length time profile:

Name	Name of the profile.
Description	Description of the profile (optional).
GPS coordinates	GPS coordinates indicating the physical location of the camera(s) assigned to the profile.
Sunrise offset	Number of minutes (+/-) by which activation of the profile is offset by sunrise.
Sunset offset	Number of minutes (+/-) by which deactivation of the profile is offset by sunset.
Time zone	Time zone indicating the physical location of the camera(s).

Manage notification profiles

With notification profiles you can set up ready-made e-mail notifications, which can automatically be triggered by a rule (see "Manage rules" on page 165), for example when a particular event occurs. You can even include still images and AVI video clips in the email notifications.

Note that when using the SMTP Service with .NET 4.0, it is not possible to send attachments over 3 MB. However, two hotfixes (must be installed on the management server in the listed order) from Microsoft® can be found at:
http://connect.microsoft.com/VisualStudio/Downloads/DownloadDetails.aspx?DownloadID=30226
http://connect.microsoft.com/VisualStudio/Downloads/DownloadDetails.aspx?DownloadID=31723

TLS (Transport Layer Security) and its predecessor SSL (Secure Socket Layer) is not supported; if the sender belongs on a server that requires TLS or SSL, e-mail notifications will not work properly. Also, you may be required to disable any e-mail scanners that could prevent the application sending the e-mail notifications.

Prerequisites

Before you can create notification profiles, you must specify settings for the outgoing SMTP mail server you are going to use for the e-mail notifications.

Optionally, if you want the notification profile's e-mail notifications to be able to contain AVI video clips, the compression settings for use when generating the AVI files must also be specified.

1. Go to the Management Client's menu bar, and select *Tools* > *Options...* This will open the *Options* window.
   - For **outgoing SMTP Mail Server**: Specify settings for the outgoing SMTP mail server (see "Outgoing SMTP mail server settings" on page 209) on the *Mail Server* tab.
   - For **AVI Compression**: Specify AVI compression settings (see "Specify AVI compression settings" on page 209) on the *AVI Generation* tab.

Add notification profiles

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Rules and Events*, right-click *Notification Profiles*, and select *Add Notification Profile...* This will open the *Add Notification Profile* wizard.
2. On the wizard's first step, specify name and description.

   Click *Next*.
3. On the wizard's second step, verify that *Email* is selected, click *Next*.
4. On the wizard's third step, specify recipient, subject, message text and time between e-mails:

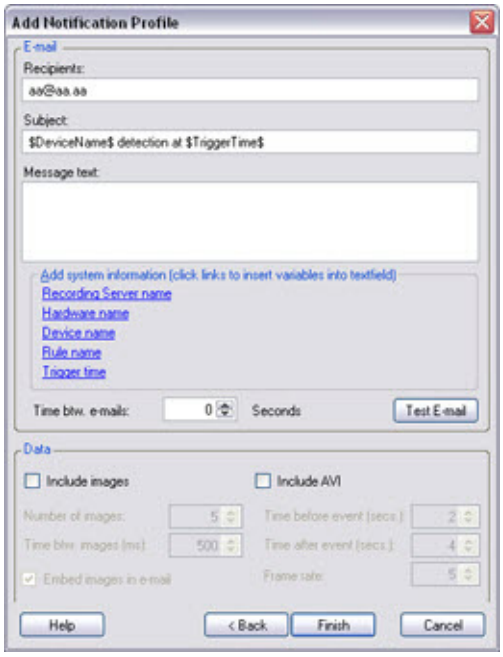

5. If you want send a test e-mail notification to the specified recipients, click *Test E-mail*.
6. If you want to include pre-alarm still images in e-mail notifications under the notification profile, select *Include images*, and specify number of images, time between images and whether images should be embedded in e-mail or not.
7. If you want to include AVI video clips in e-mail notifications under the notification profile, select *Include images*, and specify time before and after event and frame rate.
8. Click *Finish*.

Use rules to trigger e-mail notifications

You use the *Manage Rule* for creating rules. The wizard takes you through all required steps. You specify the use of a notification profile during the step on which you specify the rule's actions:

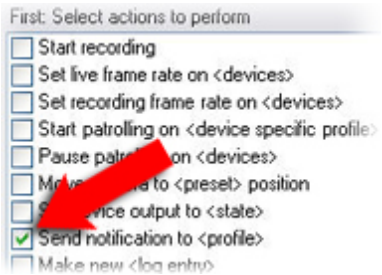

When selecting the action *Send notification to <profile>*, you get the option of selecting the required notification profile. You also get the option of selecting which cameras any recordings to be included in the notification profile's e-mail notifications should come from:

Example only; in *Manage Rule*, you click the links to make your selections

Keep in mind that recordings cannot be included in the notification profile's e-mail notifications unless something is actually being recorded.

If still images or AVI video clips are required in the notification profile's e-mail notifications, you should therefore verify that the rule you are creating— or another existing rule— specifies that recording should take place. The following example is from a rule which includes both a *Start recording* action and a *Send notification to ...* action:

For more information about rules in general, refer to Manage rules (on page 165).

Notification profile settings

Name	Description
Name	Type a descriptive name for the notification profile. The name appears later whenever you select the notification profile during the process of creating a rule.

Name	Description
Description (optional)	Type a description of the notification profile. The description appears when you pause your mouse pointer over the notification profile in the Overview pane (see "Panes Overview" on page 33)'s *Notification Profiles* list.
Recipients	Type the e-mail addresses to which the notification profile's e-mail notifications should be sent. To type more than one e-mail address, separate addresses with a semicolon. Example: aa@aaaa.aa;bb@bbbb.bb;cc@cccc.cc
Subject	Type the text you want to appear as the subject of the e-mail notification. You can insert system variables, such as *Device name*, in the subject and message text field. To insert variables, click the required variable links in the box below the field.
Message text	Type the text you want to appear in the body of the e-mail notifications. In addition to the message text, the body of each e-mail notification automatically contains this information: ▸ What triggered the e-mail notification. ▸ The source of any attached still images or AVI video clips
Time between e-mail	Specify required minimum time (in seconds) to pass between the sending of each e-mail notification. Examples: ▸ If specifying a value of *120*, a minimum of 2 minutes will pass between the sending of each e-mail notification, even if the notification profile is triggered again by a rule before the 2 minutes have passed. ▸ If specifying a value of *0*, e-mail notifications will be sent each time the notification profile is triggered by a rule. This can potentially result in a very large number of e-mail notifications being sent. If using the value *0*, you should therefore carefully consider whether you want to use the notification profile in rules which are likely to be triggered frequently.
Number of images	Specify the maximum number of still images you want to include in each of the notification profile's e-mail notifications. Default is five images.
Time between images (ms)	Specify the number of milliseconds you want between the recordings presented on the included images. Example: With the default value of 500 milliseconds, the included images will show recordings with half a second between them.
Embed images in e-mail	If selected (default), images will be inserted in the body of e-mail notifications. If not, images will be included in e-mail notifications as attached files.
Time before event (secs.)	This setting is used to specify the start of the AVI file. By default, the AVI file will contain recordings from 2 seconds before the notification profile is triggered. You can change this to the number of seconds you require.
Time after event (secs.)	This setting is used to specify the end of the AVI file. By default, the AVI file will end 4 seconds after the notification profile is triggered. You can change this to the number of seconds you require.
Frame rate	Specify the number of frames per second you want the AVI file to contain. Default is five frames per second. The higher the frame rate, the higher the image quality and AVI file size.

Manage user-defined events

If the event you require is not on the *Events Overview* list, you can create your own user-defined events. Such user-defined events can be useful if you want to integrate other systems with your surveillance system.

> **Example**: With user-defined events, you can use data received from a third-party access control system as events in the system; the events can subsequently trigger actions. This way, you can, for example, begin recording video from relevant cameras when somebody enters a building.

User-defined events can also be used for manually triggering events while viewing live video in the Ocularis Client or automatically if used in rules (see "Manage rules" on page 165).

> **Example:** When user-defined event 37 occurs, PTZ camera 224 should stop patrolling and go to preset position 18.

Through roles (see "Specify rights of a role" on page 187), you define which of your users should be able to trigger the user-defined events.

User-defined events can be used in two ways, simultaneously if required:

- **For providing the ability to manually trigger events in the Ocularis Client**

  In this case, user-defined events make it possible for end users to manually trigger events while viewing live video in the Ocularis Client. So, when a user-defined event occurs because an Ocularis Client user triggers it manually, a rule can trigger that one or more actions should take place on the system.

- **For providing the ability to trigger events through API**

  In this case, user-defined events can be triggered from outside the surveillance system. Using user-defined events this way requires that a separate API (Application Program Interface; a set of building blocks for creating or customizing software applications) is used when triggering the user-defined event. Authentication through Active Directory is required for using user-defined events this way. This ensures that even though the user-defined events can be triggered from outside the surveillance system, only authorized users will be able to do it.

  Also, user-defined events can via API be associated with meta-data, defining certain devices or device groups. This is highly usable when using user-defined events to trigger rules: you avoid having a rule for each device, basically doing the same thing. Example: A company uses access control, having 35 entrances, each with an access control device. When an access control device is activated, a user-defined event is triggered in the system. This user-defined event is used in a rule to start recording on a camera associated with the activated access control device. It is defined in the meta-data which camera is associated with what rule. This way the company does not need to have 35 user-defined events and 35 rules triggered by the user-defined events; a single user-defined event and a single rule are enough.

  When user-defined events are used this way, you may not always want them to be available for manual triggering in the Ocularis Client.

Whichever way you choose to use user-defined events, each user-defined event must first be added through the Management Client.

Note that if you rename a user-defined event, the Management Server must be Refreshed on the Ocularis Base.

Also note that if you delete a user-defined event, this will affect any rules in which the user-defined event is used. Furthermore, a deleted user-defined event will not disappear from Ocularis Clients immediately; only after the Management Server is Refreshed on the Ocularis Base.

Add a user-defined event

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Rules and Events*, and select *User-defined Events*.
2. In the Overview pane (see "Panes Overview" on page 33), right click *Events* and select *Add User-defined Event...*

3. Type a name for the new user-defined event, and click *OK*. The newly added user-defined event will now appear in the list in the Overview pane.
4. If the user has rights to do so (refer to About roles (on page 182)), the user-defined event can now be manually triggered from Ocularis Client. Remember to create one or more rules (see "Manage rules" on page 165) specifying what should take place when the custom event occurs.

Rename a user-defined event

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Rules and Events*, and select *User-defined Events*.
2. In the Overview pane (see "Panes Overview" on page 33), select the required user-defined event.
3. In the Properties pane (see "Panes Overview" on page 33), overwrite the existing name.
4. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

Security

About security

In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), you can work with the following under *Security*:

- ***Roles*: (see "About roles" on page 182)** Roles determine which of your system's features users and groups (see "Manage users and groups" on page 182) can use. In other words, roles determine rights and handles security within the application.
- ***Basic Users*** (see "***About basic users***" on page 192): Basic users are much like Windows users—but specific to a OnSSI Federated Architecture site (see "About OnSSI Federated Architecture" on page 212).

About roles

Depending on the recording component, functionality described here may be limited or unavailable.

When you work with roles, you must first create the role, then add some users/groups, and, if relevant, a time profile (see "Manage time profiles" on page 173). One role is predefined in the system, and cannot be deleted: the *Administrators* role. In addition to the *Administrators* role, you can add as many roles as required in your organization.

To manage roles in the system, go to the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Security*, and select *Roles*.

For more information see:

- Manage users and groups (on page 182)
- Assign and remove users and groups to/from roles
- Manage roles (on page 184)
- Specify rights of a role (on page 187).

Manage users and groups

With Ocularis, you need only one user on the Management Server to be used to import the system into Ocularis Base. The additional content here regarding users and roles is for informational purposes only.

In the system, you define roles (see "About roles" on page 182) first, then you add users/groups to the roles. Roles determine which of the system's features users and groups are able to use. In other words, roles determine rights.

Once you have defined roles, you can add users and groups. See Assign and remove users & groups to/from a role.

PREREQUISITES

A server with Active Directory installed that acts as domain controller must be available on your network before you can add users and groups through the Active Directory service. Consult your network administrator if in doubt.

ADD USERS AND GROUPS THROUGH ACTIVE DIRECTORY (NORMAL WAY)

Users and groups are normally added from Active Directory, although users can also be added without Active Directory.

Using Active Directory for adding existing user and group information to the system has several benefits: the fact that users as well as groups are specified centrally in Active Directory means that you do not have to create any user accounts from scratch in the system. It also means that you do not have to configure any authentication of users on the system. Authentication is handled by Active Directory.

What is Active Directory? Active Directory is a distributed directory service included with several Windows Server operating systems. It identifies resources on a network in order for users or applications to access them. Users as well as groups are specified centrally in Active Directory.

ACTIVE DIRECTORY USER AND GROUP CONCEPTS

Active Directory uses the concepts of users and groups.

Users

Users are Active Directory objects representing individuals with a user account. Example:

Groups

Groups are Active Directory objects capable of containing several users. In this example, the Management Group has three members (i.e. it contains three users):

Groups can contain any number of users. By adding a group to the system, you add all of its members in one go. Once you have added the group to the system, any changes made to the group in Active Directory (such as new members added or old members removed) at a later stage are immediately reflected in the system.

Note that a user can be a member of more than one group at a time.

ADD USERS NOT USING ACTIVE DIRECTORY

While you primarily add users and groups to roles (see "About roles" on page 182) through Active Directory, you can also add individual users—but not groups—without Active Directory. If you do not use Active Directory, note the following:

- When you install the management server, the user under which the Management Server service runs must be a local PC user on the server.
- On the computer running the management server, simple file sharing must be disabled the following way:
  1. On the computer running management server, right-click *Start*, and select *Explore*.
  2. In the window that opens, select the *Tools* menu, then select *Folder Options...*
  3. Select the *View* tab.
  4. Scroll to the bottom of the *Advanced settings* list, and make sure that the *Use simple file sharing (Recommended)* check box is cleared.
  5. Click OK, and close the window.

- You add users to roles through the Management Client almost as when adding users from Active Directory. However, when adding users, you must refer to particular users on particular computers, as in this example where the user *USER001* on the computer *PC001* is added:

When users added this way log in to the system, the user must *not* specify any server name, PC name, or IP address as part of their user names. Example of a correctly specified user name: USER001. Example of an incorrectly specified user name: *PC001*/*USER001*. The users should of course still specify their passwords, etc.

Manage roles

Roles determine which of your system's features users and groups (see "Manage users and groups" on page 182) are able to use. In other words, roles determine rights and handles security within the application.

You define roles first. Added roles automatically also become view groups.

One role is predefined in the system, and cannot be deleted: the *Administrators* Role.

In addition to the *Administrators* role, you are able to add as many roles as required in your organization.

To manage roles in the system, expand the *Security* folder in the Management Client's Site Navigation pane (see "Panes Overview" on page 33), and select *Roles*.

For more information refer to Assign and remove users and groups to/from roles and Specify rights of a role **(on page 187).**

Roles may also determine access to views in clients.

Note that renaming a role will not change the name of a view group based upon the role.

MORE ABOUT ADMINISTRATORS ROLE

The *Administrators* role is predefined, and cannot be deleted. Users and groups with the *Administrators* role have complete and unrestricted access to the entire system. For this reason it is not necessary to specify role settings for the *Administrators* role.

You add users and groups to the *Administrators* role just as with any other role; refer to Assign and remove users and groups to/from roles.

Role settings tabs are not available for *Administrators* role as users and groups with this role have unrestricted access to the system

IMPORTANT: Users with *local machine administrator* rights on the computer running the management server will automatically have administrator rights on the management server. It is therefore important that you verify which users have *local machine administrator* rights on the computer running the management server: Only users whom you trust as administrators of your system should have *local machine administrator* rights on the computer running the management server.

ADD A ROLE

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Security*, and right-click *Roles*.
2. Select *Add Role*. This will open the *Add Role* dialog.
3. Type a name and description of the new role:

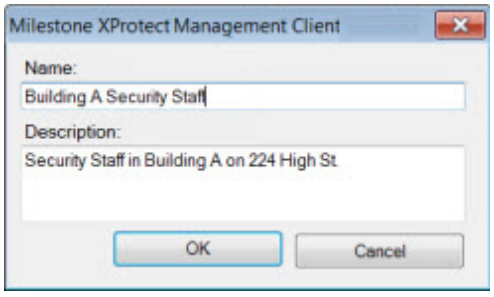

4. Click *OK*.
5. The new role is added to the *Roles* list in the Overview pane (see "Panes Overview" on page 33). By default, a new role does not have any users/groups associated with it.
6. You are now able to assign users/groups to the role, and to specify which of the system's features they should be able to access. Refer to Assign and remove users and groups to/from a role and Specify rights of a role (on page 187).

COPY A ROLE

If you have a role with complicated settings and/or rights and need a similar—or almost similar—role, it might be easier to copy the already existing role and make minor adjustments to the copy than to creating a new role from scratch.

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Security*, click *Roles,* right-click the required role in the Overview pane (see "Panes Overview" on page 33), select *Copy Role...*
2. In the dialog that opens, give the copied role a new unique name and description.
3. Click *OK*.

DELETE A ROLE

Before deleting a role (see "About roles" on page 182), keep in mind that you are able to delete a role even when users and/or groups have been assigned to the role. It is therefore often a good idea to verify if any users/groups are assigned to the role before deleting it.

Verify if any users/groups are assigned to a role

- In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Security*, and right-click *Roles*.
- Select the required role in the Overview pane (see "Panes Overview" on page 33), then select the *Users and Groups* tab in the Properties pane (see "Panes Overview" on page 33). Any users and/or groups assigned to the role will be listed on the *Users and Groups* tab.

Delete a role

1. In the Management Client's Site Navigation pane, expand *Security*, and right-click *Roles*.
2. Right-click the unwanted role in the Overview pane, and select *Delete Role*.
3. Click *Yes*.

If you delete a role, this does not delete a view group based upon the role. For information about deleting view groups, refer to Manage view groups (on page 135).

VIEW EFFECTIVE ROLES

With the Effective Roles feature, you are able to view all roles (see "About roles" on page 182) of a selected user or group (see "Manage users and groups" on page 182). This ability is especially convenient if you are using groups; in fact it is the only way of viewing the roles of individual group members.

1. Open the *Effective Roles* window. There are three ways in which you can open the *Effective Roles* window:
   - From the Management Client's menu bar, by selecting *Tools* > *Effective Roles...*
   - From the Overview pane (see "Panes Overview" on page 33) (when working with roles), by right-clicking anywhere inside the pane, then selecting *Effective Roles*...
   - From the Site Navigation pane (see "Panes Overview" on page 33), by expanding *Security*, then right-clicking *Roles*, then selecting *Effective Roles*...
2. In the *Effective Roles* window's *User name* field, type the user name of the required user.

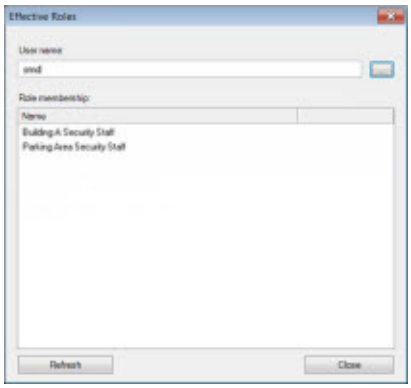

3. If you typed the user name directly into the *User name*, click *Refresh* in the lower part of the window to display the roles of the user.

   If you used Active Directory to browse for the user, the user's roles will be displayed automatically.

Work with users, groups and roles

To assign or remove Windows users or groups or basic users (see "About basic users" on page 192) to/from a role, do the following:

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand ***Security***, and select ***Roles***. Then select the required role in the Overview pane (see "Panes Overview" on page 33):

2. In the Properties pane (see "Panes Overview" on page 33), select the ***Users & Groups*** tab at the bottom.
3. Click ***Add...***, select between ***Windows user*** or ***Basic user***.

ASSIGN WINDOWS USERS AND GROUPS TO ROLE

1. Select Windows user. This opens the ***Select Users***, ***Computers and Groups*** dialog:

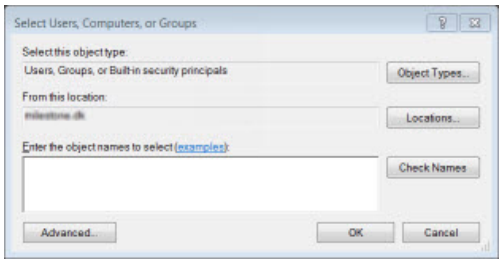

2. Verify that the required object type is specified. If, for example, you need to add a computer, click ***Object Types*** and mark ***Compute***r. Also verify that the required domain is specified in the ***From this location*** field. If not, click ***Locations...*** to browse for the required domain.
3. In the ***Enter the object names to select*** box, type the required user names, initials, or other types of identifier which Active Directory can recognize.

   Tip: Typing part of a name is often enough. Use the Check Names feature to verify that the names, initials, etc. you have typed are recognized by Active Directory.
4. Click ***OK***. The selected users/groups are now added to the ***Users & Groups*** tab's list of users who have been assigned the selected role.

ASSIGN BASIC USERS TO ROLE

a Select ***Basic User***. This opens the ***Select Basic Users*** to add to ***Role*** dialog:

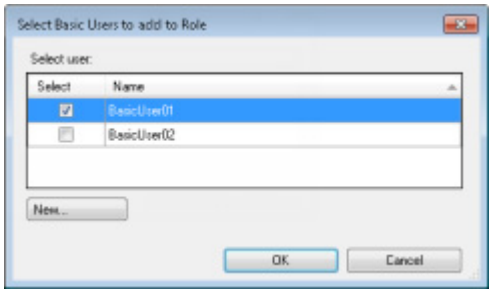

b Select the basic user(s) that you want to assign to this role.

Optional: Click ***New...*** to create a new basic user.

c Click ***OK***. The selected basic user(s) are now added to the ***Users & Groups*** tab's list of basic users who have been assigned the selected role.

REMOVE USERS AND GROUPS FROM ROLE

Bear in mind that a user may also have roles through group memberships. When that is the case, you cannot remove the individual user from the role. Group members may also hold roles as individuals. To find out which roles users, groups, or individual group members have, use the Effective Roles (see "Manage roles" on page 184) feature.

a On the ***Users & Groups*** tab, select the user or group you want to remove, then click ***Remove*** in the lower part of the tab.

Tip: You can select more than one user or group, or a combination of groups and individual users, if required.

b Confirm that you want to remove the selected user(s) or and group(s). Click ***Yes***.

Specify rights of a role

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Security*, and select *Roles*. Then select the required role in the Overview pane (see "Panes Overview" on page 33):

2. In the Properties pane (see "Panes Overview" on page 33), specify required rights for the role on the relevant tabs:

The next item/section is only relevant if you run Ocularis ES.

USERS AND GROUPS

RIGHTS

Refer to Assign/remove users and groups to/from roles.

DEVICE RIGHTS

The *Device* tab lets you specify which features users/groups with the selected role should be able to use for each device (e.g. a camera) or device group in the Ocularis Client.

The tab is divided into two halves: In the upper half, select the item for which you want to specify role rights. In the lower half, specify which role right should apply for users/groups with the selected role in the Ocularis Client.

Remember to repeat for each required device/device group. You can also select a device group, and specify role rights for the entire device group in one go.

Camera-related rights

Depending on the recording component, functionality described here may be limited or unavailable.

Setting determines whether...

- ***AUX Commands :***.it will be possible to use auxiliary commands from the Ocularis Client.

  What are AUX Commands? AUX is short for Auxiliary. Such commands offer the user control of, for example, wipers on a camera connected via a video server. Camera-associated devices connected via auxiliary connections are controlled from the Ocularis Client.

- ***Add :*** it will be possible to add bookmarks in recorded video from the Ocularis Client.
- ***Delete :*** it will be possible to delete bookmarks in recorded video from the Ocularis Client.
- ***Edit :*** it will be possible to edit bookmarks in recorded video from the Ocularis Client.
- ***View :*** It will be possible to view bookmarks in recorded video from the Ocularis Client.
- ***View live within time profile*** : live viewing of video from the selected camera(s) will be possible in access clients.
- ***Export video :*** .the database export feature can be used when browsing recorded video from selected camera(s) in the Ocularis Client. Furthermore, the AVI, JPEG and export features can be used in similar way in all access clients.
- ***Get sequences :*** the *Sequences* feature can be used when browsing recorded video from the selected camera(s) in access clients.
- ***Playback Video :***.playing back of recorded video from the selected camera(s) will be possible in access clients.
- ***Smart Search:*** the *Smart Search* feature can be used when browsing recorded video from the selected camera(s) in the Ocularis Client.
- ***Visible :*** the selected camera(s) will be visible in access clients.

The View live *right also requires that the role has been granted the right to view the access clients'* Live *mode. This right is granted as part of the application rights.*

The Export Video *and* Playback Video *rights also require that the role has been granted the right to view the access clients'* Browse *tab. This right is granted as part of the application rights.*

Microphone-related rights

Depending on the recording component, functionality described here may be limited or unavailable.

Setting determines whether...

- ***Visible*** **:** ...the selected microphone(s) will be visible in the Ocularis Client.
- ***Listen to live audio:*** ...listening to live audio from the selected microphone(s) will be possible in the Ocularis Client.
- ***Browse audio:*** ...browsing of recorded audio from the selected microphone(s) will be possible in the Ocularis Client.
- ***Export audio*** **:** ...the export feature can be used when browsing recorded audio from the selected microphone(s) in the Ocularis Client.
- ***Get sequences:*** **This feature is currently not supported** ...the *Sequences* feature can be used when browsing recorded audio from the selected microphone(s) in the Ocularis Client.

Speaker-related rights

Depending on the recording component, functionality described here may be limited or unavailable.

Setting determines whether...

- ***Visible*** **:** the selected speaker(s) will be visible in the Ocularis Client.
- ***Listen to live audio :***listening to live audio from the selected speaker(s) will be possible in the Ocularis Client.
- ***Browse audio :*** browsing of recorded audio from the selected speaker(s) will be possible in the Ocularis Client.
- ***Export audio :*** the export feature can be used when browsing recorded audio from the selected speaker(s) in the Ocularis Client.
- ***Get sequences*** **: This feature is currently not supported** ...the *Sequences* feature can be used when browsing recorded audio from the selected speaker(s) in the Ocularis Client.

IMPORTANT: *Although what is being said through a speaker can be recorded and archived (see "About storage and archiving" on page 61), there is currently no way of playing back or exporting such recorded outgoing audio. Therefore, some of the speaker-related rights currently have no effect. Features for playing back and exporting recorded outgoing audio, etc. will be available in subsequent releases as soon as possible.*

Input-related rights

- ***Visible*** **:** Determines whether information about the selected input(s) will be visible to users of the Ocularis Client.

Output-related rights

Depending on the recording component, functionality described here may be limited or unavailable.

Setting determines whether...

- ***Visible*** **:**.the selected output(s) will be visible in the Ocularis Client. If visible, the output will be selectable on a list in the Ocularis Client.
- ***Activate output :*** the selected output(s) can be activated from the Ocularis Client.

Outputs are selected and activated from Ocularis Client.

Why are some check boxes filled with squares? Square-filled check boxes can only appear if you are specifying role rights for a device group, in which case they indicate that the right in question currently applies for some, but not all, devices within the device group.

Square-filled check boxes indicate that settings currently apply for some, but not all, devices within a device group

You can still select or clear such square-filled check boxes, but note that your choice will in that case apply for *all* devices within the device group. Alternatively, select the individual devices in the device group to verify exactly which devices the right in question applies for.

PTZ RIGHTS

Depending on the recording component, functionality described here may be limited or unavailable.

Relevant only if PTZ (Pan/Tilt/Zoom) cameras are available on your system, the *PTZ* tab lets you specify which features users/groups with the selected role should be able to use in the Ocularis Client.

The tab is divided into two halves: In the upper half, select the item for which you want to specify role rights. In the lower half, specify which role right should apply for users/groups with the selected role in the Ocularis Client. Note that only PTZ cameras and device groups containing PTZ cameras are available for selection

The following rights are available:

- ***Allow PTZ Control:*** Determines if the selected role is able to use the pan, tilt and zoom features of the selected PTZ camera(s).
  - False: Users/groups with the selected role will not be able to use the pan, tilt and zoom features of the selected PTZ camera(s)
  - True: Users/groups with the selected role will be able to use the pan, tilt and zoom features of the selected PTZ camera(s)
- ***PTZ Priority:*** Determines the priority of PTZ cameras. When several users on a surveillance system wish to control the same PTZ camera at the same time, conflicts may occur. This setting lets you alleviate the problem by specifying a priority for use of the selected PTZ camera(s) by users/groups with the selected role. Specify a priority from *1* to *32.000*, where *1* is the lowest priority.

  Default PTZ priority is 3000.

  Example: You specify that the role *Security Manager* should have very high priority when using a PTZ camera, whereas the role *Security Assistant* should have low priority when using the PTZ camera. Now, if a user with the role *Security Manager* and a user with the role *Security Assistant* want to control the PTZ camera at the same time, the user with the role *Security Manager* will win the ability to control the camera.

  If your system is upgraded from an older version of the system, the old values (*Very Low, Low, Medium, High* and *Very High*) have been translated as follows:
  - *Very Low = 1000*
  - *Low = 2000*
  - *Medium = 3000*

- *High = 4000*
- *Very High = 5000*

Users of the Ocularis Client are able to stop/resume a patrolling PTZ camera's patrolling. This PTZ feature is not regulated by PTZ priority.

- ***Allow activation of PTZ presets:*** Determines if the selected role is able to move the selected PTZ camera(s) to preset positions.
  - False: Users/groups with the selected role will not be able to move the selected PTZ camera(s) to preset positions
  - True: Users/groups with the selected role will be able to move the selected PTZ camera(s) to preset positions

SPEECH RIGHTS

Depending on the recording component, functionality described here may be limited or unavailable.

Relevant only if loudspeakers are available on your system.

The tab is divided into two halves: In the upper half, select the item for which you want to specify role rights. In the lower half, specify which role right should apply for users/groups with the selected role in the Ocularis Client.

The following rights are available:

- ***Speak live:*** Determines whether users with the selected role will be able talk through the selected speaker(s).
- ***Speak priority:*** When several Ocularis Client users want to talk through the same speaker at the same time, conflicts may occur. This setting lets you alleviate the problem by specifying a priority for use of the selected speaker(s) by users/groups with the selected role. Specify a priority from *Very low* to *Very high*.

  Example: You specify that the role *Security Manager* should have very high priority when talking through a speaker, whereas the role *Security Assistant* should have low priority when talking through the speaker. Now, if a user with the role *Security Manager* and a user with the role *Security Assistant* want to talk through the speaker at the same time, the user with the role *Security Manager* will win the ability to talk.

If two users with the same role want to speak at the same time, the first-come first-served principle applies.

For the right to work, the role must also be granted the right to view the Ocularis Client in *Live* mode. This right is granted as part of the application rights. Furthermore, the speaker(s) must be *visible* in Ocularis Client; something you determine as part of the device rights.

APPLICATION RIGHTS

Depending on the recording component, functionality described here may be limited or unavailable.

The *Application* tab lets you specify client access to specific functions. First select a specific time profile or ***Always.*** Next, select the functions that the role should have access to***:***

- *Status API:* N/A.
- *Service Registration API:* N/A.
- ***Reports***: The configuration report (see "About configuration report" on page 195) functionality within ***System Dashboard***.

EXTERNAL EVENT RIGHTS

Depending on the recording component, functionality described here may be limited or unavailable.

The tab is divided into two halves: In the upper half, select the item for which you want to specify role rights. In the lower half, specify which role right should apply for users/groups with the selected role in the Ocularis Client.

The following rights are available:

- ***Trigger external event with time profile:*** In the Ocularis Client it is possible to manually trigger your surveillance system's external events.

VIEW GROUP RIGHTS

The tab is divided into two halves: In the upper half, select the item for which you want to specify role rights. In the lower half, specify which role right should apply for users/groups with the selected role in the Ocularis Client.

The following rights are available:

- ***Visible*** **:** Determines if the selected role is able to see the selected view group (and any views contained in the view group) in clients.
- ***Modify*** **:** Determines if the selected role is able to make changes to the selected view group (and any views contained in the view group) in clients.
- ***Delete*** **:** Determines if the selected role is able to delete the selected view group (and any views contained in the view group) in clients.
- ***Create subgroups and views*** **:** Determines if the selected role is able to create subgroups and views in the selected view group.

REMOTE RECORDING RIGHTS

Depending on the recording component, functionality described here may be limited or unavailable.

The tab is divided into two halves: In the upper half, select the item for which you want to specify role rights. In the lower half, specify which role right should apply for users/groups with the selected role in the Ocularis Client.

The following rights are available:

- ***Retrieve remote recordings***: Determines if users/groups with the selected role should be able to retrieve remote recordings (see "Remote recording - camera/remote system" on page 128).

SERVERS RIGHTS

The next item/section is only relevant if you run Ocularis ES.

Specifying role rights on the *Servers* tab is only relevant if you have integrated Ocularis CS servers into your system; refer to Manage Ocularis CS servers (see "Manage Ocularis CS servers" on page 202) for more information.

NETMATRIX RIGHTS

Specifying role rights on the *NetMatrix* tab is only relevant if you have configured NetMatrix recipients on your system.

From the Ocularis Client Limited Mode it is possible to send video to selected NetMatrix recipients. The *NetMatrix* tab lets you specify which NetMatrix recipients should be selectable for this purpose in the Ocularis Client (running in Limited Mode).

The tab is divided into two halves: In the upper half, select the item for which you want to specify role rights. In the lower half, specify which role right should apply for users/groups with the selected role in the Ocularis Client.

The following rights are available:

- ***Visible*** **:** Determines whether users/groups with the selected role will be able to select and send video to the NetMatrix recipient from the Ocularis Client (Limited Mode).

About basic users

When working with basic users, it is important to understand the difference between basic user and Windows user.

- Basic users are authenticated by a user name/password combination and are specific to a OnSSI Federated Architecture site (see "About OnSSI Federated Architecture" on page 212). Even if basic users have

the same name and password, a basic user created at one federated site does not have access to another federated site.

- Windows users are authenticated based on their Windows login and are specific to a machine.

Manage basic users

Once you have created a basic user, you must add it to a role (and add this role to a group), if you want to use it actively in your system. Refer to About roles (on page 182) for details.

ADD BASIC USER

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand the ***Security*** node, right-click ***Basic Users***, and select *Create Basic Users*. This opens the *New Basic Authentication User* dialog.
2. Fill in the needed properties (see "Basic user properties" on page 193).
3. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

BASIC USER PROPERTIES

- ***User name:*** Name of basic user.
- ***Description:*** Description of basic user (optional).
- ***Password:*** Enter user name.
- ***Repeat password:*** Re-enter user name.

System dashboard

About system dashboard

In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), you are able to work with the following under *System Dashboard*:

- ***System Monitor:*** (see "About system monitor" on page 194)Here you can view and print detailed system reports on servers, devices and cameras.
- ***Current Task:*** (see "About current task" on page 195)Here you can get an overview of tasks under a selected recording server.
- ***Configuration Report:*** (see "About configuration report" on page 195)From here you can decide what to include in your system configuration reports—and print them.

About system monitor

From the Site Navigation pane (see "Panes Overview" on page 33), expand *System Dashboard*, and click *System Monitor*. This brings up the System Monitor using embedded browser technology.

If you access the system monitor from a **server** operating system, you might experience a message regarding *Internet Explorer Enhanced Security Configuration*. Follow instructions in the message in order to add the System Monitor page to the *Trusted sites zone* before proceeding.

A Data Collector Server service is dedicated to collect performance counter values on servers and cameras to be used in the System Monitor functionality.

About Data Collector Server service

The Data Collector Server service is automatically installed on the same machines as the management, recording and log server(s).

Normally, the Data Collector Server service requires no maintenance. However, if the service ***does*** stops, it will result in missing live feed to the System Monitor (clearly indicated in the system monitor by error texts). On the machine where the Data Collector Server service is installed, do the following to restart it:

1. In Windows' *Start* menu, select *Control Panel*, and then...
    - If using *Category* view, find the *System and Security* category and click *Administrative Tools*.
    - If using *Small icons* or *Large icons,* click *Administrative Tools*.
2. Double-click ***Services***.
3. Locate the ***OnSSI Data Collector Server***. Right-click it. From the menu that appears, select ***Start*** to restart the service.

Work with system monitor

Use the <, > and home icons to navigate the System Monitor.

From here you can view system information and create reports on:

- ***Management server:*** shows data on *your management server*
- ***Recording servers:*** shows data on *any number of recording servers* in your surveillance setup, which can be viewed per:
    - ***Disks***
    - ***Storage***

- ***Network***
- ***Cameras***

- ***Failover recording servers:*** shows data on *any number of failover recording servers* in your surveillance setup
- ***Additional servers:*** shows data on *log servers, etc*. in your surveillance setup
- ***Cameras:*** shows data *on any camera in any camera group* in your surveillance setup.

Each of these corresponds to a clickable, expandable area, most of which contains sub-areas. Each sub-area represents a server. When clicked, they provide relevant dynamic data on this server.

The *Cameras* bar however, contains a list of camera-groups to select from. Once a group is selected, you can select a specific camera and see dynamic data for it.

All servers display **CPU usage** and **available memory** information. Furthermore, recording servers also display **connection status** information.

Within each view, you can find a *History* link. Click it to view historic data and reports (to view reports on a camera, click the name of the camera). For each historic report, you can view data for the last 24 hours, 7 days or 30 days.

If you want to save and/or print reports, click the *Send to PDF* icon.

About current task

To get an overview of tasks under a selected recording server, their begin time, estimated end time and progress, do the following:

From the Site Navigation pane (see "Panes Overview" on page 33), expand *System Dashboard*, and click *Current Task*.

In general, all information showed in *Current Tasks* are snapshots and are refreshed by clicking on the refresh button in the lower right corner of the Properties pane (see "Panes Overview" on page 33).

About configuration report

When creating your pdf configuration reports, you can include any possible elements of your system which you want to see in the report. Examples of what can be included ranges from licenses over device to alarm configuration, and much more.

Furthermore, you can customize your font and page setup and include a customized front page as listed:

Add a configuration report

1. From the Site Navigation pane (see "Panes Overview" on page 33), expand *System Dashboard* and click *Configuration Reports*. This brings up the report configuration page.
2. Select the elements that you want to include in your report.
3. Optional: Click *Front Page...* to customize your front page. In the window that appears, fill in the needed info.

   Remember to select Front page *as an element to include in you report, otherwise the front page you customize will not be included in your report.*
4. Click *Formatting...* to customize your font, page size and margins. In the window that appears, select the wanted settings.
5. When you are ready to export, click *Export...* and select a name and save location for you report.

Tip: Remember, not all fonts support all special characters. If you have trouble viewing your special characters, try selecting a different font.

Configure report details

The following is available when setting up reports:

- ***Select All:*** Selects all elements in the list
- ***Clear All:*** Clears all elements in the list
- ***Front Page:*** Opens a dialog allowing you to customize the front page
- ***Formatting:*** Opens a dialog allowing you to format the report
- ***Export:*** Opens a dialog allowing you to select the save location for the report and create the pdf.

Server logs

Manage logs

In the Management Client, you are able to view and copy contents from different logs related to the management server. The different logs have different purposes:

- Audit Log records user activity.
- Event Log records event-related information (see "Events overview" on page 161).
- Rule Log records rules (see "Manage rules" on page 165) in which the *Make new <log entry>*action (see "Actions and Stop actions" on page 136) has been specified.
- System Log records system-related information.

Your system has a number of default settings related to the different logs, refer to Handle log settings (on page 200). Furthermore, you are able to view logs in a number of different languages, export them, and save the exported logs as tab delimited text (.txt) files at a location of your choice; refer to Export log (on page 199).

View log

To view a log, expand the *Management Server Logs* item in the Management Client's Site Navigation pane (see "Panes Overview" on page 33), then select appropriate the log.

Read and copy logs

Each row in a log represents a log entry. A log entry contains a number of information fields which are listed and briefly explained. Note, it is also possible to double-click any row and have all its details presented in a Log Details window. From the Log Details window, it is also possible to copy/paste any log contents:

- ***Level***
  - **All logs:** Display an icon indicating the level of the log entry:

    indicates info
    indicates error
    indicates warning.
- ***UTC Time***
  - **All logs *:*** Timestamp in coordinated universal time (UTC), an international high-precision time standard.
- ***Local Time***
  - **All logs:** Timestamp in the local time of your system's server.
- ***Description***
  - **All logs:** Description of the logged incident.
- ***Source Type***
  - **Rule Log only:** Type of equipment on which the logged incident occurred. Since log entries are administrator-defined and relate to incidents in your system, source type will normally be *System*.
  - **Event and System Logs only:** Type of equipment on which the logged incident occurred, for example *Management Server* or *Device*.
  - **Audit Log only:** Type of equipment on which the logged incident occurred. Since remote user access is handled by the management server, source type will typically be *Server*.
- ***ID***
  - **All logs:** Identification number of the logged incident.

- ***Event Type***
  - **All logs, except Audit Log:** Type of event represented by the logged incident.

    For more information about event types, refer also to the events overview (on page 161).
- ***Source Name***
  - **All logs:** Name of the management server, device, etc. on which the logged incident occurred.
- ***Service Name***
  - **Event and Rule Logs only:** Name of service on which the logged incident occurred.
- ***Audit Type***
  - **Audit Log only:** Type of logged incident.
- ***Granted***
  - **Audit Log only:** Information about whether the remote user action was allowed (granted) or not.
- ***User***
  - **Audit Log only:** User name of the remote user causing the logged incident.
- ***Location***
  - **Audit Log only:** IP address or host name of the computer from which the remote user caused the logged incident.
- ***Rule Name***
  - **Rule Log only:** Name of the rule triggering the log entry.
- ***Generator Type***
  - **Rule Log only:** Type of equipment on which the logged incident was generated. Since the log entries are administrator-defined and relate to incidents in your system, generator type will normally be *System*.
- ***Generator Name***
  - **Rule Log only:** Name (if any) of the equipment on which the logged incident was generated.

Navigate log

If a log contains more than one page of information, you are able to navigate between the log's pages by clicking the buttons in the bottom right corner of the log pane:

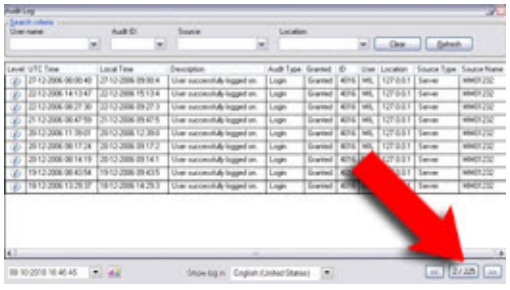

<< lets you move one step towards the log page containing the most recent log entries.

1 / 171 indicates which page you are currently viewing (e.g. page 1 of 171). By clicking the button, you are able specify a page number and go straight to that page.

>> lets you move one step towards the log page containing the oldest log entries.

Furthermore, 08-10-2010 10:46:45 in the lower left corner lets you jump to a specific date and time in the log.

Change log language

1. In the bottom part of the log pane, in the *Show log in* drop down-box, select the wanted language.

2. The log is displayed in the selected language.

Next time you open the log, it is reset to the default language.

Search log

To search a log, use the *Search criteria* box in the top part of the log pane:

1. Specify your search criteria by selecting the required user name, location, etc. from the lists.

   Tip: You can combine selections, or make no selection in certain lists, as required. The more search criteria you combine, the less search results you will typically get.

2. Click *Refresh* to make the log page reflect your search criteria.

   Tip: To clear your search criteria, and return to viewing all of the log's content, click *Clear*.

Export log

You are able to export logs, and save the exported logs as tab delimited text (.txt) files at a location of your choice.

Example of an exported log .txt file

Example of an exported log .txt file viewed in Notepad.

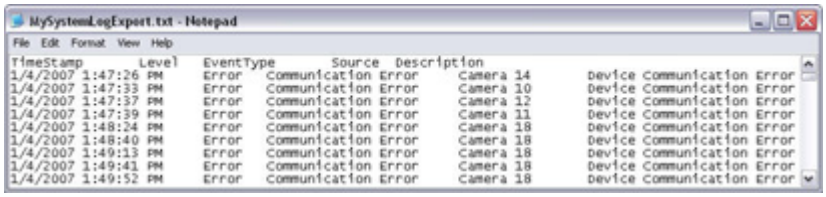

You are able to target the exported log content by specifying which log, which log elements, and which time range to include in the export. For example, you are able to specify that only the System Log's error-related log messages from between January 2nd 2007 08:00:00 and January 4th 2007 07:59:59 should be included in your export.

To export a log, do the following:

1. Expand the *Management Server Logs* item in the Management Client's Site Navigation pane (see "Panes Overview" on page 33), and select the required log.
2. If you want to target the exported log's content, select the required criteria in the *Search criteria* section above the log. For example, you may select that your export should only contain log messages at a particular level, such as errors or warnings.

   Remember to click *Refresh* to make the log page reflect your selected criteria.

3. In the Management Client's menu bar, select *Action* > *Export Log...* This will open the *Export Log* window:

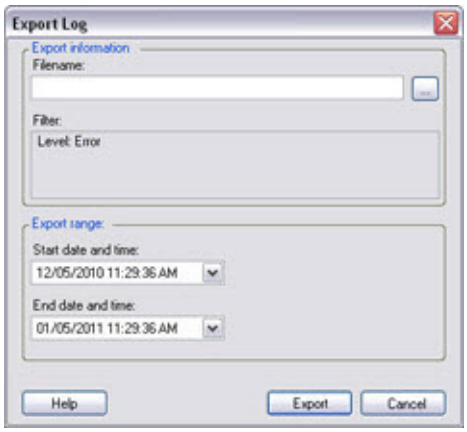

4. In the *Export Log* window's *Filename* field, specify a name for the exported log file.

   By default, exported log files will be saved in your *My Documents* folder. However, you are able to specify a different location by clicking the browse button next to the field.

5. Any criteria you have selected in order to target the content of the exported log will be listed in the *Filters* field. The field is non-editable; if you find that you need to change your criteria, close the window, and repeat steps 2-4.

6. Specify the time period you want the export to cover. You do this by specifying the required boundaries in the *Start date and time* and *End date and time* fields respectively. By clicking the arrow, you are able to select the required date from a calendar:

   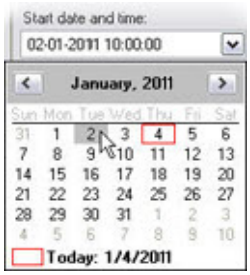

   To specify an exact time, overwrite the required time elements (hours:minutes:seconds) with the required values. In this example, the hours element is being overwritten:

   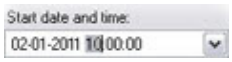

   Tip: When you have selected time elements in order to overwrite them, you are also able to use your keyboard's UP ARROW and DOWN ARROW keys to increase/reduce the numbers in increments of one unit.

7. Click *Export* to export the required log content to the required location.

Handle log settings

Your system has a number of default settings related to its logs. To verify or change these settings, do the following:

1. Go to the Management Client's menu bar, and select *Tools* > *Options...*
2. In the *Options* window, select the *General* tab. One of the tab's settings applies for all types of logs:
   - ***Number of log rows to retrieve per page*:** Lets you specify the number of log rows you want to view on a single log page. If a log contains more than the specified number of rows, you will be able view the remaining rows on subsequent log pages.
3. Go to the *Options* window's *Management Server Logs* tab:

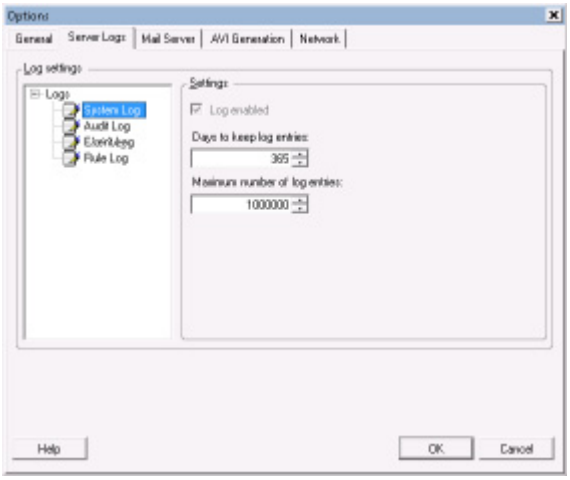

In the tab's left box, select the required log. The selected log's settings are displayed in the tab's right box:

- ***Log enabled:*** Lets you enable/disable the selected log. By default, all logs—**except the Event Log**—are enabled.

  The *System Log* and *Audit Log* cannot be disabled by clearing the box.

- ***Days to keep log entries:*** Lets you specify how many days the log's information should be kept for. Default is 30 days.

  Excess log content will be deleted if the log reaches its maximum allowed size (see *Maximum number of entries*) before the specified number of days is reached.

- ***Maximum number of entries*** **:** Lets you specify the maximum size of the log. Default is 50.000 entries.

  Excess log content will be deleted if it reaches its maximum allowed age (see *Days to keep log entries*) before the specified number of entries is reached.

For the *Audit Log*, you will also see:

- ***Enable user access logging*** **:** Lets you include detailed information about specific user actions in the audit log, e.g. about users' viewing of live video (and associated audio), PTZ actions, activation of output and events, export, playback of video and audio, use of playback features, any denied access to features, etc.

- ***Playback sequence logging length:*** Lets you specify how long a playback sequence may last and still be considered and logged as **one** sequence.

  Example: If you select 60 seconds, you may view 60 consecutive seconds of playback video but still only leave one log entry in the Audit Log. Specifying a high number of seconds may help limit the number of viewed sequences logged, and, in this way, reduce the size of the audit log.

- ***Records seen before logging:*** Lets you specify the number of records to be viewed before logging the sequence.

4. Click *OK*.

Ocularis CS

Manage Ocularis CS servers

This section is only relevant if you run Ocularis ES.

If your organization has Ocularis CS installations, you can integrate Ocularis CS servers into your system. You do this by adding the Ocularis CS servers through the Management Client.

This integration only works with Ocularis CS servers. For the current version of Ocularis LS, use OnSSI Federated Architecture (see "About OnSSI Federated Architecture" on page 212) for adding Ocularis LS servers as children. Furthermore, integration is also not possible if your system uses IPv6.

When added, Ocularis CS servers can send data and video to your system. You can compare added Ocularis CS servers with recording servers and these will likewise be available for viewing in clients.

Note that roles defined in the Management Client can be given access to data from Ocularis CS servers. This is done by coupling roles in your system with Ocularis CS user rights.

Furthermore, Ocularis CS servers added in the Management Client will be listed in the *Add/Remove Ocularis CS Servers* dialog which you can open by selecting *Ocularis CS Servers...* from the *Tools* menu.

Ocularis CS's Recording Server service must be running for your system to receive data from the Ocularis CS installation.

Limitations when adding Ocularis CS servers

When using Ocularis, it is not necessary to add RC-C servers to the RC-E system as this is handled at the Ocularis level using the Ocularis Administrator application. The information here is for informational purposes only and for legacy systems.

There are a few limitations to how Ocularis CS servers will work when added as slaves to your system. They will provide operational status and status details on cameras and Ocularis CS servers but not on any other device types.

Also, you cannot define cameras, user rights, scheduling, or other settings for the Ocularis CS installation, or see previews of the cameras in your system. All necessary Ocularis CS settings must be made in Ocularis CS's *Administrator* application or other relevant Ocularis CS applications.

For client users, it will be completely transparent whether feeds come from an Ocularis CS server or from a recording server in your system. The users have access to cameras depending on their roles defined in the Management Client. If a role has borrowed user rights from an added Ocularis CS server, users with that role have access to data from the Ocularis CS server according to the borrowed user rights.

Prerequisites for access roles for Ocularis CS servers

On the Ocularis CS server, open the *Image Server Administrator* window to see if one of the Ocularis CS users has user rights that can be used in connection with a role in your system.

Write the Ocularis CS user's user name and password or Windows account down. You will need this information when you use the Management Client to define roles with access to Ocularis CS servers. Note that user names and passwords are case sensitive.

You can also create a new user in Ocularis CS, and assign the required user rights in Ocularis CS, so they match the role in your system. Refer to the Ocularis CS documentation for more information about creating new users in Ocularis CS.

Define roles with access to Ocularis CS servers

To give access to data from Ocularis CS servers, do the following in the Management Client:

1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Security,* and select *Roles*.

2. Select the required role from the list. If you want to define a new role, refer to About roles (on page 182) for more information.
3. At the bottom of the *Role Settings* pane select the *Servers* tab.
4. Select the Ocularis CS server to which you want to assign the role.
5. Select a user with the Ocularis CS user rights that represent the correct user rights for the role (in your system) you are assigning it to. You can do this in two ways:
   - In the *Basic Authentication* section, enter the user name and password for a user which is defined as basic authenticated user in Ocularis CS.

     or -
   - In the *Windows Authentication* section, enter the Windows account name for a user which is defined as a Windows authenticated user in Ocularis CS.

   Tip: If in doubt whether a user is defined as a Basic or Windows authenticated user in Ocularis CS, open the *Image Server Administrator* window on the Ocularis CS server, and click *User Setup....* Refer to the Ocularis CS documentation for more information

The selected Ocularis CS user has not automatically been assigned to the role in question through the Management Client. The user's Ocularis CS user rights have just been borrowed by the role, but the actual user has not been assigned to the role.
The system does not verify that the specified user name or password is correct or that the specified user name, password or Windows account name correspond to a defined user in Ocularis CS. Therefore, make sure that you enter the information correctly. Note also that user names and passwords are case sensitive.
6. In the toolbar (see "Management Client Overview" on page 30), click *Save*.

Add Ocularis CS servers

To add an existing Ocularis CS installation to your system, do the following:

1. From the Management Client's *Tools* menu select Ocularis CS *Servers...*
2. In the *Add/Remove Ocularis CS Servers* dialog click *Add...*.
3. Enter the IP address or the host name of the required Ocularis CS server in the *Ocularis CS server IP / Host name* field.
4. Enter the port number used by the Ocularis CS server's Image Server in the *Port number* field.

   Tip: The default port number is 80; if in doubt, you can find the port number in the *Image Server Administrator* window on the Ocularis CS server.
5. Now enter information about the administrator of the Ocularis CS server. You can do this in two ways:
   - Select *Windows* and click the browse button to the right of the *User name* field to use the Windows authentication method which authenticates the administrator through the administrator's Windows login.

   - or -
   - Select *Basic* and enter the Ocularis CS administrator's user name and password in the *User name* and *Password* fields.

The reason why it is important that you enter the Ocularis CS administrator information, is that you as administrator then will have unlimited rights to data from both your system and the Ocularis CS installation.

The connection to the Ocularis CS server is now established, but no roles in the Management Client—except the Administrator role—have been given access to data from the Ocularis CS server. Refer to Define roles with access to Ocularis CS servers (see "Define roles with access to Ocularis CS servers" on page 202) for more information about giving users access to data from added Ocularis CS servers.

Remember to define the network configuration settings, so the management server will be able to handle the token authentication of clients for added Ocularis CS servers.

In the Management Client, you must add all Ocularis CS servers you would like to receive data from. The Ocularis CS system's internal master/slave setup cannot be reused by your system.

Ocularis CS server network configuration

With the network configuration settings you specify the management server's server address so that the management server can handle the token authentication of clients for added Ocularis CS servers.

From the Management Client's *Tools* menu select Ocularis CS *Servers...*

1. In the *Add/Remove Registered Services* window, click *Network...*
2. Specify the LAN and/or WAN IP address of the management server.

   If all involved servers (both the management server and the trusted servers or the required Ocularis CS) are on your local network, you can simply specify the LAN address. If one or more involved servers access the system through an internet connection, you must also specify the WAN address.

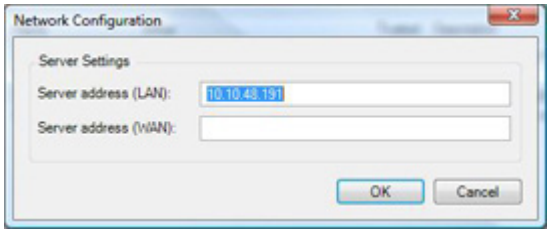

3. Click *OK*.

Edit Ocularis CS servers

From the Management Client's *Tools* menu select Ocularis CS *Servers...*

1. Select an Ocularis CS server from the list, and click *Edit...* in the *Add/Remove Ocularis CS Servers* dialog.
2. Edit the relevant settings and click *OK*.

Registered services

Manage registered services

Occasionally, you have servers and/or services which should be able to communicate with the system even though they are not directly part of the system. Some services, but not all, can register themselves automatically in the system. Services that can automatically be registered are:

- Log Server service (see "Management Server" on page 8)

Automatically registered services are displayed in the list of registered services.

You can manually specify servers/services as registered services in the Management Client:

Access registered services configuration

1. In the Management Client's menu bar, select *Tools* > *Registered Services...*
2. The *Add/Remove Registered Services* window opens. From this window you can manage registered services.

Add and edit registered services

1. In the *Add/Remove Registered Services* window, click *Add...* or *Edit...*, depending on your needs.
2. In the *Add Registered Service* or *Edit Registered Service* window (depending on your earlier selection), specify or edit settings (see "Registered services settings" on page 205).
3. Click *OK*.

Manage network configuration

With the network configuration settings you specify the management server's server LAN and WAN addresses in order for the management server and the trusted servers to be able to communicate.

1. In the *Add/Remove Registered Services* window, click *Network...*
2. Specify the LAN and/or WAN IP address of the management server.

   If all involved servers (both the management server and the trusted servers or the required Ocularis CS) are on your local network, you can simply specify the LAN address. If one or more involved servers access the system through an internet connection, you must also specify the WAN address.

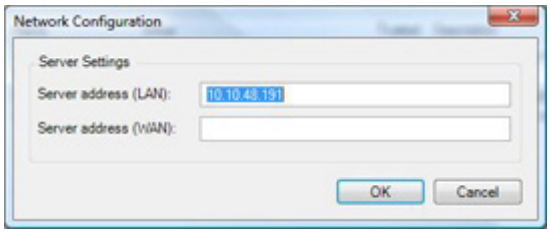

3. Click *OK*.

Registered services settings

In the *Add Registered Service* or *Edit Registered Service* window, specify the following:

Name	Description
Services type	Prefilled field.
Name	Name of the registered service. The name is only used for display purposes in the Management Client.
Description	Description of the registered service. The description is only used for display purposes in the Management Client.
URL	Click *Add* to add the IP address or hostname of the registered service in question. If specifying a hostname as part of a URL, the host in question must exist and be available on the network. URLs must begin with *http://* or *https://* and must not contain any of the following characters: *< > & ' " * ? | [ ] "*. **Example** of a typical URL format: http://ipaddress:port/directory (where port and directory are optional). Note that you can add more than one URL if required.
External	Select if the registered service connects to the management server with a public IP address.
Trusted	Select if the registered service should be trusted immediately (this is often the case, but the option gives you the flexibility to add the registered service and then mark it as trusted by editing the registered service later). Note that changing the trusted state will also change the state of other registered services sharing one or more of the URLs defined for the relevant registered service.

Options

Options

The Management Client's *Options* dialog lets you specify a number of settings related to the appearance of the application, to logging, to mail server configuration, etc.

You access the *Options* dialog from the Management Client's menu bar (see "Management Client Overview" on page 30), by selecting *Tools* > *Options*.

The *Options* dialog features the following tabs:

General

Specify the following from ***Options***' *(see "*Options*" on page 207)General* tab:

Name	Description
Number of log rows to retrieve per page	Select how many rows to appear on a single log page. The default value is 50 rows. If a log contains more rows than your selected value, it displays the additional rows on additional pages.
Default preview frame rate	Select which frame rate to use for the thumbnail camera images displayed in the Preview pane (see "Panes Overview" on page 33). Default is 1 frame per second. Refreshing the Management Client' layout (by pressing F5 on your keyboard or selecting *Action* > *Refresh* from the menu bar) is required for a change to take effect. Note that a high frame rate (that is, a high image quality) in combination with a large number of thumbnail images in the Preview pane may slow the system down. You can limit the number of thumbnail images with the *Max. number of previews* setting.
Max no. of previews	Select the maximum number of thumbnail images displayed in the Preview pane. Default is 64 thumbnail images. Refresh the Management Client's layout to make changes take effect. Note that a large number of thumbnail images in combination with a high frame rate (that is, a high image quality) may slow the system down. You can limit the frame rate used for the thumbnail images with the *Default preview frame rate* setting.
Motion detection 'on' when adding camera devices	Select whether to enable motion detection while cameras are being added to a recording server through the *Add Hardware* (on page 53) wizard. Select the check box to enable motion detection while using the wizard (default). Note that this setting only applies while *Add Hardware* is in use. When the wizard is not in use, motion detection is active for all cameras for which it has been enabled, regardless of this setting.
Enable multicast live when adding camera devices	Select whether to enable multicast (see "Multicasting tab (recording server properties)" on page 75) while cameras are being added to a recording server through the wizard Add Hardware. Select the check box to enable multicast while using the wizard (default). Note that this setting only applies while *Add Hardware* is in use. When the wizard is not in use, multicast will be active for all cameras for which it has been enabled, regardless of this setting.

Language	Select the language of the Management Client. Available languages: English, French, Spanish. Restart the Management Client to make language changes take effect.
Timeout for PTZ sessions	Ocularis Client users with necessary user right can manually interrupt the handling of PTZ cameras. This setting lets you select how much time should pass before regular patrolling is resumed after a manual interruption. The setting applies for all PTZ cameras on your system.
Ignore device communication errors if communication reestablished before	Select how long a communication error may last without the system log logging it or, in other words, when it is brief enough to be ignored.

Tip: Motion detection is a key element in the surveillance system, and is by default enabled for all cameras on the system. However, motion detection uses a relatively large amount of computing resources. If your system features a very large number of cameras, and you have enabled motion detection on all cameras, the system may be slowed down slightly, and adding of new cameras may take longer than usual. To add new cameras as quickly as possible, you have the option of disabling motion detection while the wizard *Add Hardware* is in use.

Server logs

The ***Options***' *(see "*Options*" on page 207)Server Log* tab lets you specify settings for the system's five different management server logs.

See Manage logs (on page 197) for more information.

Mail server

The ***Options (on page 207)'*** *Mail Server* tab lets you specify settings for the outgoing SMTP mail server you are going to use with your system:

Name	Description
Sender e-mail address	Type the e-mail address you want to appear as the sender of e-mail notifications for all notification profiles. Example: sender@organization.org.
Outgoing mail (SMTP) server name	Type the name of the SMTP mail server which will be used for sending e-mail notifications for all notification profiles. Example: mailserver.organization.org.

AVI generation

The ***Options***' (see "Options" on page 207) *AVI Generation* tab lets you specify compression settings for the generation of AVI video clip files. Specifying these settings is necessary if you want to include AVI files in e-mail notifications sent out by rule-triggered notification profiles (see "Manage notification profiles" on page 176).

Specify the following from ***Options (on page 207)'****AVI generation* tab:

Name	Description
Compressor	Select the required codec (compression/decompression technology). Indeo® 5.10 (if available) generally provides a good compromise between quality and file size. You can configure some, but not all codecs.

Compression quality	(Not available for all codecs). Use the slider to select the required degree of compression (*0-100*) to be performed by the codec. *0* means no compression, generally resulting in high image quality and large file size. *100* means maximum compression, generally resulting in low image quality and small file size. If the slider is not available, compression quality is determined entirely by the selected codec.
Keyframe every	(Not available for all codecs). If you want to use keyframes, select the check box and specify the required number of seconds between keyframes in the neighboring field. A keyframe is a single frame stored at specified intervals. The keyframe contains the entire view of the camera, whereas the following frames contain only the pixels that change. This helps greatly reduce the size of files. If the check box is not available, or not selected, every frame contains the entire view of the camera.
Data rate	(Not available for all codecs). If you want to use a particular data rate, select the check box and specify the required number of kilobytes per second in the neighboring field. If the check box is not available, or not selected, date rate is determined entirely by the selected codec.

Network

The ***Options***' *(see "*Options*" on page 207)Network* tab lets you specify local IP address ranges.

Refer to Manage local IP address ranges (on page 210) for more information.

User settings

The ***Options***' (see "Options" on page 207) Use*r Settings* tab lets you specify settings for user preference, such as whether a message should be shown when edge recording is enabled.

Refer to *Record* tab overview (see "Record tab overview" on page 125) for more information.

Specify AVI compression settings

Outgoing SMTP mail server settings

When you configure outgoing SMTP mail server settings, specify the following:

Name	Description
Sender e-mail address	Type the e-mail address you want to appear as the sender of e-mail notifications for all notification profiles. Example: sender@organization.org.
Outgoing mail (SMTP) server name	Type the name of the SMTP mail server which will be used for sending e-mail notifications for all notification profiles. Example: mailserver.organization.org.

AVI compression settings

When you set up AVI compression settings, specify the following:

Name	Description
Compressor	Select the required codec (compression/decompression technology). Indeo® 5.10 (if available) generally provides a good compromise between quality and file size. You can configure some, but not all codecs.
Compression quality	(Not available for all codecs). Use the slider to select the required degree of compression (*0-100*) to be performed by the codec. *0* means no compression, generally resulting in high image quality and large file size. *100* means maximum compression, generally resulting in low image quality and small file size. If the slider is not available, compression quality is determined entirely by the selected codec.
Keyframe every	(Not available for all codecs). If you want to use keyframes, select the check box and specify the required number of seconds between keyframes in the neighboring field. A keyframe is a single frame stored at specified intervals. The keyframe contains the entire view of the camera, whereas the following frames contain only the pixels that change. This helps greatly reduce the size of files. If the check box is not available, or not selected, every frame contains the entire view of the camera.
Data rate	(Not available for all codecs). If you want to use a particular data rate, select the check box and specify the required number of kilobytes per second in the neighboring field. If the check box is not available, or not selected, date rate is determined entirely by the selected codec.

Manage local IP address ranges

When the Ocularis Client connects to a surveillance system, an amount of initial data communication, including the exchange of contact addresses goes on in the background. This happens automatically, and is completely transparent to users.

Clients may connect from the local network as well as from the internet, and in each case the surveillance system should be able to provide suitable addresses so the clients can get access to live and recorded video from the recording servers:

When clients connect locally, the surveillance system should reply with local addresses and port numbers.

- When clients connect from the internet, the surveillance system should reply with the recording servers' public addresses (see "Network tab (recording server properties)" on page 78), i.e. the address of the firewall or NAT (Network Address Translation) router, and often also a different port number (which is then forwarded to recording servers).

The surveillance system must therefore be able to determine whether a client belongs on a local IP range or on the internet. For this purpose, you can define a list of IP ranges which the surveillance system should recognize as coming from a local network.

Working with Local IP Address ranges...

1. In the Management Client's menu bar, select *Tools* > *Options*. This will open the Options dialog (see "Options" on page 207).

Tip: You can also access the *Options* dialog from the *Network* tab; this can be handy if you are also configuring the public IP address of a recording server.

2. In the *Options* dialog, select the *Network* tab.

Define local IP address ranges

On the *Network* tab, click *Add*.

a In the *Range Start* column, specify the first IP address in the required range. Then specify the last IP address in the range in the *Range End* column.

Tip: If required, a range may include only one IP address (example: 192.168.10.1-192.168.10.1).

b If more ranges are required, repeat steps a - b.

c Click *OK*.

Edit local IP address ranges

a Overwrite the existing information in the *Range Start* and *Range End* columns as required.

b Click *OK*.

OnSSI Federated Architecture

About OnSSI Federated Architecture

This section is only relevant if you run Ocularis ES.

OnSSI Federated Architecture allows multiple individual standard systems (also known as sites) to interconnect in a parent/child hierarchy of sites.

Federated Architecture works with Ocularis ES and Ocularis LS servers.

IMPORTANT: Federated hierarchy is only possible with version 4.0 or newer of Ocularis ES. Before installing the system, refer to Important prerequisites when running federated sites (on page 212).

In this text, the term *parent* refers to a parent site and *child* to a child site.

Through Federated Architecture, client users—based on their user rights—have seamless access to video, audio and other resources across individual sites. In addition, through a single login, administrators can centrally manage all sites within the federated hierarchy—again based on administration rights for the individual sites.

As it provides unlimited scalability, flexibility and accessibility to video surveillance across multiple sites and has no limit to the number of sites you can add, Federated Architecture is well suited for large installations covering multiple buildings, campuses, or entire city areas.

Each site in a federated hierarchy is installed and configured as a normal stand-alone system with standard system components, settings, rules, schedules, administrators, users, and user rights. Once each site has been installed, these can be connected by requesting a Federated Architecture link from one site (the parent) to another (the child). When the link is established, the two sites automatically create a Federated Architecture hierarchy to which more sites can be added to grow the federated hierarchy.

Illustration of OnSSI Federated Architecture (on page 218)

In this example, the federated hierarchy consists of six sites. As illustrated, each site can be both a parent and a child at the same time making it possible to create a hierarchy with as many levels as needed. It is also evident that a site can link to several child sites on the same level in a hierarchy.

Once a federated hierarchy is created, it allows users and administrators logged in to a site, to access that site and any child or sub-child sites it may have. Access to child and sub-child sites in the hierarchy is not gained automatically, but dependent on appropriate user and administrator rights.

It is only relevant to speak of a parent/child setup for management servers (see "Management Server" on page 8)—not for recording servers. However, due to their relations to management servers, recording servers will automatically become part of the parent/child setup.

Refer to Manage OnSSI Federated Architecture (on page 219) for details on how to work with federated architecture.

Important prerequisites when running federated sites

The easiest way to make federated architecture work correctly is to prepare your system for this feature during installation. There are certain important prerequisites that you must ensure already at the time of installing your management server. This can be done in different ways - choose between the procedures in alternative 1-3:

Alternative 1: Connect sites from the same domain (with common domain user) and customize the installation of the management server to federated architecture

Before installation of the management server, a common domain user should be created and used as the administrator on all computers involved in the MFA. Depending on whether you select *Custom* or *Typical* during installation of the management server, make sure to select the appropriate procedure. Note that a typical installation requires more configuration on all sites before federated architecture will work properly.

Custom installation:

1. Start the management server installation (see "Installation overview" on page 13) and select *Custom*.
2. Select to install the Management Server service using a user account.

   The selected user account must be the administrator on all management servers and must also be used when installing the other management servers in the federated architecture setup.
3. Finish the installation.
4. Repeat steps 1-3 to install any other systems you want to connect in the federated architecture.
5. Refer to Add site to hierarchy (on page 220) for details on how to proceed with the federated architecture.

Typical Installation - set up network service on all servers:

1. Start the management server installation (see "Installation overview" on page 13) and select *Typical*, let it run till it finishes.

   This will install the management server as a network service.
2. Repeat step 1 to install any other systems you want to connect with the federated architecture.
3. Using a Management Client, connect to the management server you want to have as your parent site.
4. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand *Security*, click *Roles,* and in the Overview pane (see "Panes Overview" on page 33), click *Administrator*.
5. Add the child computer to this parent server's *Administrator* role. Refer to Assign/remove users and groups to/from roles for details.
6. Log out of the parent management server and connect to the management server that you just added as a child.
7. Once again, in the Overview pane, click *Administrator*.
8. Add the parent computer to this servers *Administrator* role. Refer to Assign/remove users and groups to/from roles for details.
9. Log out of the management server, connect to the parent management server, and refer to Manage OnSSI Federated Architecture (on page 219) for details on how to proceed with federated architecture.

Alternative 2: Connect sites from different domains

To make it possible to connect sites across domains, it is very important that these domains are trusted by each other. Setting up domains to trust each other has nothing to do with federated architecture but is entirely a matter of Microsoft® Windows® Domain configuration.

For further information on how to set up trusted domains, see Microsoft® website.

1. When the domains, on which the sites you want to connect to each other in a federated architecture, are trusted correctly, follow the same instructions as if only one domain was present (see Alternative 1).

Alternative 3. Connect sites in workgroup(s)

When you connect sites inside workgroups, it is an important prerequisite for federated architecture to work correctly that the same administrator account is present on all computers you want connected in the federated architecture. This must be in place before installing the system.

1. Log in to *Windows* using a common administrator account.
2. Start the management server installation (see "Installation overview" on page 13) and click *Custom*.
3. Select to install the Management Server service using a common administrator account.
4. Finish the installation.
5. Repeat steps 1-4 to install any other systems you want to connect. They must all be installed using a common administrator account.

6. Refer to Add site to hierarchy (on page 220) for details on how to proceed.

*It is **not** possible to mix domain(s) and workgroup(s), i.e. connect sites from a domain to sites from a workgroup and vice versa.*

Licensing of OnSSI Federated Architecture

To learn about licensing in general, refer to About Licenses (on page 47).

Federated architecture can be used - freely - within the **same legal entity** as many times as needed. In a federated setup, all sites share the same software license code (see "Manage Software License Codes" on page 48) (SLC) and device licenses are shared between all sites.

In the case of **different legal entities** running federated architecture, each system requires a valid set of base and device licenses. Furthermore, in order for a device to be accessible across a federated setup, one *OnSSI Federated Architecture Device License* is required per device accessed in the federated site.

To get additional licenses for your system, contact your product dealer.

Basic rules of federated sites

- **One parent - many children**
  A child can only have one parent, but a parent can have an unlimited number of children.
- **Parent requests child, not the other way around**
  A new parent/child link is always requested by the parent, and if necessary, authorized by the child. Refer to Accept inclusion in hierarchy (on page 221).
- **One level at the time**
  A parent knows about all its children, children's children, etc., but only controls them one level down. Furthermore a child only knows about and answers to its parent one level up.
- **Synchronization of hierarchy**
  A parent always contains an updated list of all its currently attached children, children's children, etc. But when distant communication is needed, it takes place level by level, each level forwarding and returning communication, until it reaches the server requesting the information. Depending on the number of levels that must be updated, changes to a hierarchy might take a little time to become visible in the Federated Sites Hierarchy pane (see "Panes Overview" on page 33), refer to Refresh site hierarchy (on page 223). The federated hierarchy has a regularly scheduled synchronization between sites, as well as management-triggered synchronization every time a site is added or removed. This synchronization only contains site configuration data and each time will send less than 1MB. In addition to the data sent during synchronization, video or configuration data will be sent when a user or administrator views live or recorded video or configures the system. The amount of data in this case depends on what and how much is being viewed. It is not possible to schedule your own synchronizations.

Principles for setting up federated sites

When working with federated architecture, the link between management servers is established from the management server wanting to become parent to another management server. Theoretically, establishment of a parent/child relationship happens as follows:

1. The parent sends a link request to the potential child.
2. Depending on administrator settings, the child might have to authorize the link request.
3. If necessary, the child authorizes the link request.
4. Relevant info is exchanged.
5. The new parent/child link is established.

Administrators role and federated sites

- **Administrator vs. non-administrator**
  In general, you must be an administrator to work with federated architecture. However, by requesting the adding of children to a top-site (to which you have administrator rights), you can (without administrator rights to the other sites) create the overall initial infrastructure of a federation. But, as described in Manage OnSSI Federated Architecture (on page 219), the administrator of each individual child must later authorize the connection before it can take effect.
- **How to become an administrator using Active Directory - two possible scenarios**
  How to become administrator of a OnSSI Federated Architecture setup using Active Directory depends on how the management server is installed. If it is installed as described in either of the following two scenarios, you gain administrator rights of the entire setup. Otherwise not.
    - If the management server is installed as a **Network Service:** Both/All computers involved must be added as users to each other's administrator role before a parent/child link can be established without acceptance from the administrator of the child. Refer to Assign and remove users and groups to/from roles for details. This type of setup is primarily recommended if all sites in the hierarchy are not a member of the same Domain. Also refer to Important prerequisites when running federated sites (on page 212).
    - If the management server is installed as a **user account:** This user account must be a member of the administrator group of the server being linked to before one or more parent/child link(s) can be established without acceptance from the administrator of the child. This type of user right setup is primarily recommended if the number of sites in a hierarchy is large.
- **How to become an administrator using work groups**
  How to become administrator of a federated architecture setup using work groups depends on how accounts are created. If they are set up correctly, you gain administrator rights of the entire setup. Otherwise not. Refer to Important prerequisites when running federated sites (on page 212) for details on how to do this.
  If the previous criteria are not met, the administrator of a child must accept requests for inclusion in hierarchy (see "Accept inclusion in hierarchy" on page 221) manually before links can be established.
- **One or more administrators?**
  A OnSSI Federated Architecture setup can have many administrators working on it at the same time. Furthermore, the Site Navigation pane (see "Panes Overview" on page 33) is dynamic and reflects changes to the federated site made both by you and possibly other administrators. This means that you might see changes here caused by other users. You might also experience that a site you are connected to is removed from the federated site by another user. In this case, your site will be removed from the Federated Site Hierarchy pane (see "Panes Overview" on page 33), but nothing will change in the Site Navigation pane or elsewhere, allowing you to continue working.

Possibilities and constrains of federated sites

In principle, there is no limit to the number of sites you can add to federated architecture and how these can be linked, offering you unlimited scaling, flexibility and accessibility.

Frequently asked questions to federated sites

What is a federated site? A federated site is basically just an individual system, complete with management server, SQL server, one or more recording server(s), failover recording server(s) and cameras. To make use of OnSSI Federated Architecture, you must connect at least two individual systems. The Management Client is used to configure federated hierarchies. In principle, it lets you connect to any site in the federated hierarchy at any given time (if user rights permit) using the log in credentials for your home-site. This offers you a central overview, and, at the same time, lets you zoom in on selected sites by connecting to a specific site to have a closer look, make configurations, or carry out maintenance. Note however, that the Management Client is only able to see other sites from the level of the site you are logged into and downwards in the hierarchy.

What is a top-site? Your top-site is the top level management server of your entire OnSSI Federated Architecture setup.

An example an organization could have a top-level server called *MyCorp*. Second level servers called *MyCorp/RegionalServers*. Third level servers called *MyCorp/ReginalServers/CityNames.* And so on. In this case, *MyCorp* is your top-level server. There can only be one top-level server.

Tip: In a federated hierarchy, it is always a good idea to name your servers in a recognizable way, for example, using regional names or names implying where/in what context the server is located. Using, for example, consecutive numbers only, might be confusing if you have many servers.

What is a home-site? Your home-site is the site to which you are logged in. Since you may be logged in far down in the hierarchy, this is not necessarily the same as your top-site—but it **may** be. You are only able to see children from the point at which you are logged in and downwards.

Can a site be both a parent and a child at the same time? Yes, a parent with children attached, can easily be child to another site, and vice versa. This is because the parent/child concept is relative and used only in respect to other specified servers.

See the federated sites illustration (see "Illustration of OnSSI Federated Architecture" on page 218), where site 7 is the parent of site 8, but the child of site 6.

What is the difference between logging into and connecting to a site? To work with OnSSI Federated Architecture you must always be logged in to a site via the Management Client. You can log in to any site if you have administrator rights to that particular site. This is called your home-site. When logged in to your home-site, you can see all its children (if user rights permit). From your home-site you can also connect to its children (if user rights permit). Embedded in the connection process is an automated and seamless log-in, using the same credentials as your home-site log in. Connecting to a child allows you to see and work with that site (if user rights permit). However, even though technically you log out of your home-site when connecting to another site, you will still see the site structure as your (former) home-site sees it. This means, that any changes you make to a child might not be visible until such changes reaches your home-site via scheduled synchronization. So changes you make in your hierarchy might not be reflected in the Federated Sites Hierarchy pane until later.

For more details, refer to Basic rules of federated sites (on page 214).

You cannot refresh via a connection to a child, this must take place directly from the home-site.

When do I need to accept link requests? Whether as the administrator of a child you must accept a link request or not (or the link request is accepted automatically) depends on your administrator settings.

Refer to The administrator role and federated sites (see "Administrators role and federated sites" on page 215).

Where is OnSSI Federated Architecture configured and managed? Setting up and configuring OnSSI Federated Architecture takes place in the Management Client.

How do I view video from federated sites? You can view video from federated sites in any Ocularis Client, i.e. there is no need for a one-to-one relationship between sites and Ocularis Clients. You will always get the view, i.e. see the site structure as the parent you are currently logged in to.

The next item/section is only relevant if you run Ocularis ES.

Can I include Ocularis CS slave(s) in my federated hierarchy? Yes, that is possible, but only as slave(s) to a management server.

Is OnSSI Federated Architecture the same as multiple management servers, a.k.a. clustering? No, OnSSI Federated Architecture is not the same as clustering. Clustering is a method of obtaining failover support for a management server on a site. With clustering, it is only possible to have one active management server per surveillance setup, but other management servers may be set up to take over in case of failure. On the other hand, OnSSI Federated Architecture is a method of combining multiple independent sites into one large setup, offering flexibility and unlimited possibilities.

Refer to Manage OnSSI Federated Architecture (on page 219).

Federated sites example scenario—Limestone City

The following is an example of how several systems can be integrated into a federated architecture - in this case in a **City Surveillance** scenario.

Many surveillance integrators want to integrate several independent surveillance entities into a large scale system, where each site can still be used and managed locally and users and administrators can be given access to the entire large scale installation.

In this example, several governmental and business installations must be tied together in a large scale system offering the different entities local access and management of the system, as well as governmental (police etc.) access in case of crimes and emergencies.

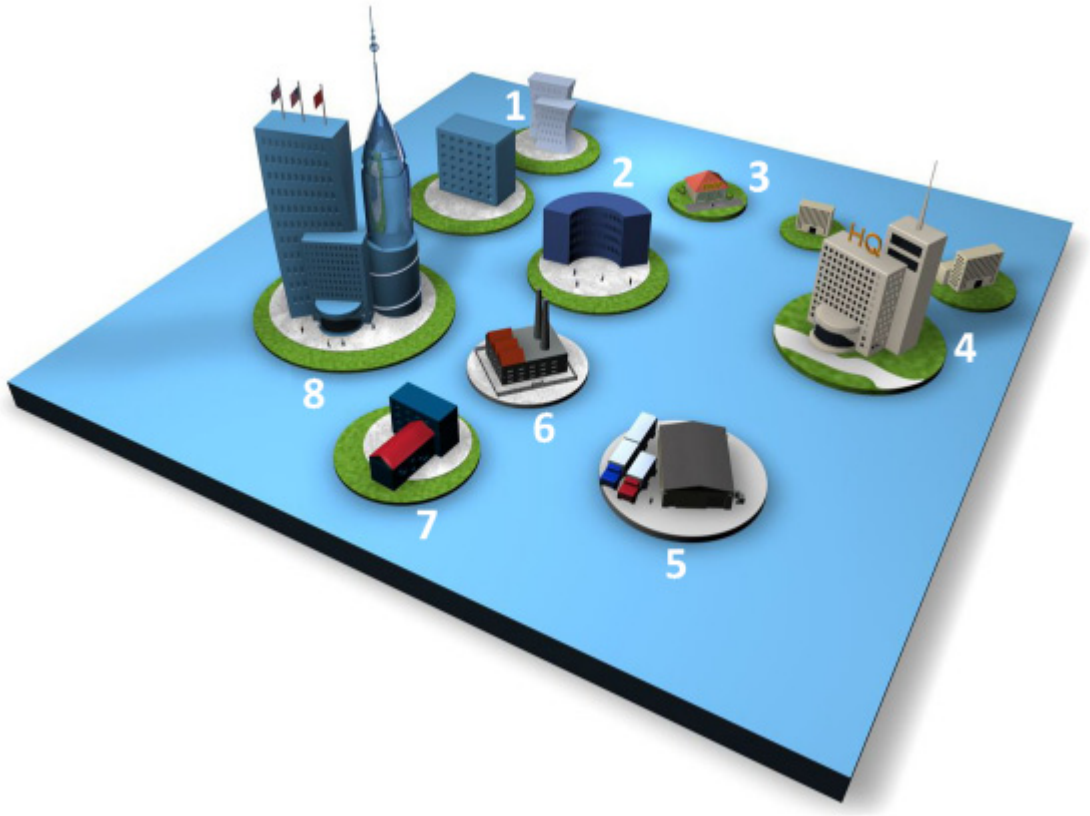

1. Downtown Residential
2. City Hall - public places
3. Residential area shops
4. A.C.M.E Industries Inc & branch offices
5. Limestone Transportation Ltd.
6. MB Industries
7. Police Headquarters
8. Limestone Center Shopping Mall

All entities must be connected to the city's video surveillance so that City Hall officials and police officers can access video from their business or residential area to monitor live video or investigate recorded video in case of break-ins, thefts, vandalism, emergencies, terror etc.

In addition to being connected to the city's video surveillance, A.C.M.E Industries Inc, Downtown Residential and Limestone Center Shopping Mall also want to segment their installation in several sites as they have several physical locations that they want to monitor. The segmented architecture offers them greater flexibility during installation and daily usage.

The city uses federated architecture, allowing the entities independent video surveillance while being tied into the city wide surveillance system at the same time.

Because the police have installations that City Hall should not have access to, the Police Headquarters is selected as the top-site in the city's federated surveillance hierarchy.

Each site is then tied into Limestone city's federated hierarchy as follows:

- **Level 1:** Police Headquarters.
    - **Level 2:** Limestone City.
        - **Level 3:** City Hall and **MB Industries as one group.**
            - **Level 4:** Central Station, Streets & Intersections and Parks as one group under City Hall.
        - **Level 3:** Limestone Center Shopping, Downtown Residential, Limestone Transportation Ltd and A.C.M.E Industries Inc. as one group.
            - **Level 4:** Shops, Branch Malls and Residential area shops as one group under Limestone Center Shopping.
            - **Level 4:** Branch Office 1 and Branch Office 2 as one group under A.C.M.E Industries Inc.

Illustration of OnSSI Federated Architecture

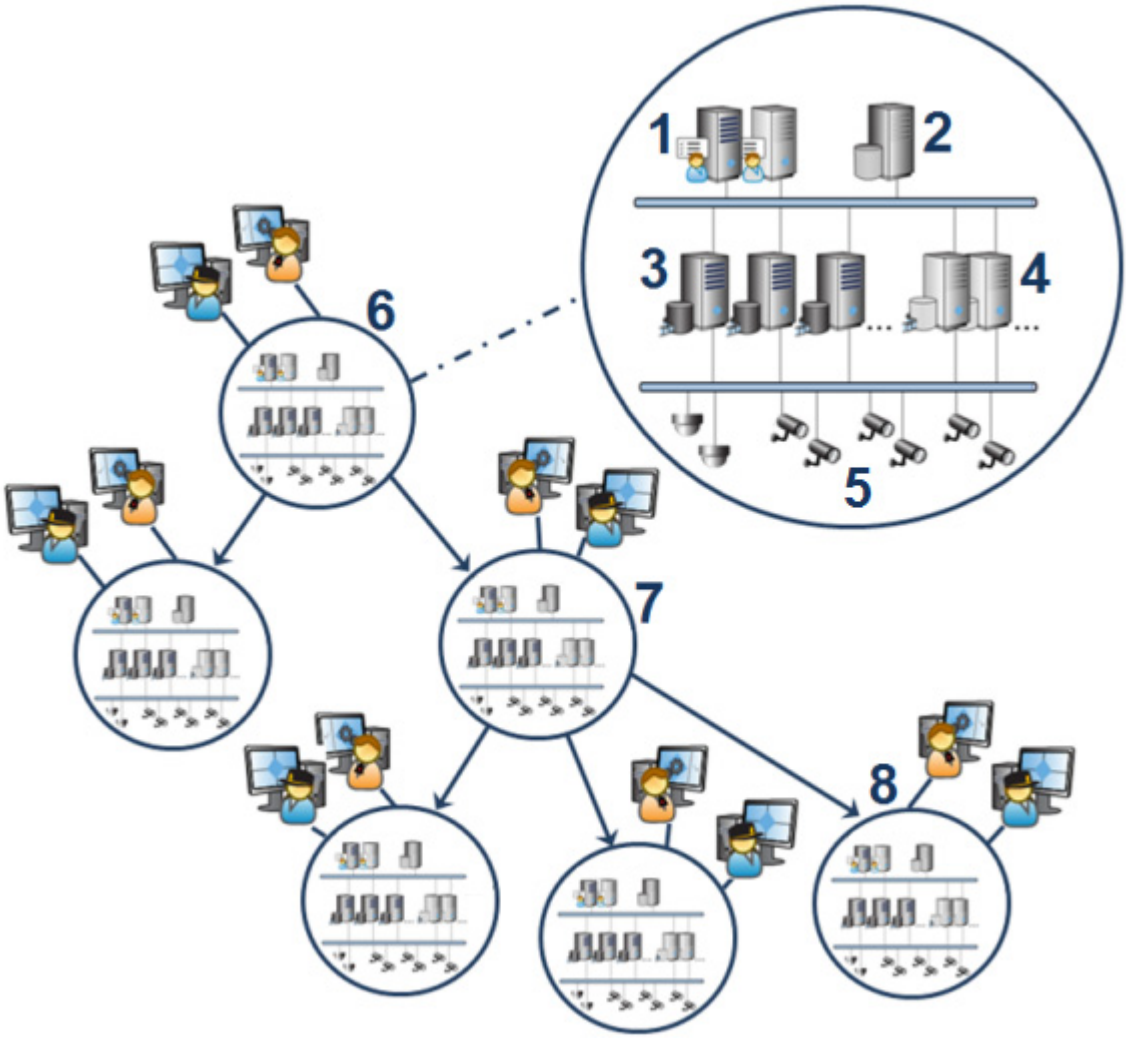

The idea behind OnSSI Federated Architecture. Parent and children linked as needed.

The contents of a federated site and parent/child setup:

1. Management server and failover recording server
2. SQL Server
3. Recording servers
4. Failover recording servers
5. Cameras
6. A Federated Site
7. Another Federated Site
8. Yet another Federated Site

Etc.

Manage OnSSI Federated Architecture

For conceptual details on OnSSI Federated Architecture, refer to OnSSI Federated Architecture overview (see "About OnSSI Federated Architecture" on page 212).

The Management Client has a Federated Sites Hierarchy pane (see "Panes Overview" on page 33) dedicated to displaying federated sites and their parent/child links. From the *View* menu (see "Management Client Menu Overview" on page 37), you can show or hide the Federated Sites Hierarchy pane. The pane is located on the left side of the Management Client window, under the Site Navigation pane (see "Panes Overview" on page 33).

The parent server you are logged in to (your home-site), is always at the top of the site hierarchy. You can view all its linked children and downwards through the parent/child hierarchy. Settings and configurations of your home-site is always reflected in the Overview and Properties panes (see "Panes Overview" on page 33) and its site-name visible at the top of the Site Navigation pane.

To connect to another site in the hierarchy (see "Connect to another site in hierarchy" on page 222), click the wanted site in the Federated Sites Hierarchy pane.

What if I only have one server and don't run federated architecture? Your user interface looks the same, but when you view the *Federated Sites Hierarchy* pane you will only see the one server in your setup.

Federated icons

There are a number of icons in federated architecture, each representing the different states a site can be in:

- Top-site in the entire hierarchy is operational:
- Top-site in the entire hierarchy is still operational but, one or more issues need attention: will be shown on top of the top-site icon.

- Normal site (not top-site) is operational:
- Normal site (not top-site) is still operational but, one or more issues need attention: will be shown on top of the normal-site icon.
- Site awaiting acceptance of inclusion in the hierarchy:
- Site being attaching, but not yet operating:

Expand/collapse

You can expand and collapse a site in the Site Navigation pane (see "Panes Overview" on page 33), to see its children, if any.

Site Navigation pane

The name, settings and configurations of the highlighted site (red arrow) are reflected in the Site Navigation pane (see "Panes Overview" on page 33).

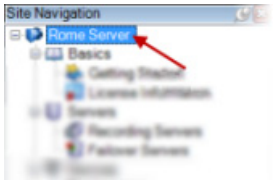

Right-click does not select

Because you must be able to delete a site without being connected to it, **right-clicking a site does not select it,** but offers a context menu, which differs depending on where in the hierarchy you are. Refer to **Action** menu (see "Management Client Menu Overview" on page 37).

Context menu

From the Federated Sites Hierarchy pane (see "Panes Overview" on page 33), a context menu lets you add sites to a hierarchy, accept inclusion in a hierarchy, rename sites in a hierarchy, detach sites from hierarchy, work with site properties and refresh site hierarchy.

Due to the nature of federated sites, when the context menu is activated from a parent, you cannot accept inclusion in the hierarchy. And when it is activated from a child, it is not possible to refresh the site hierarchy.

Add site to hierarchy

You can add children to both your home-site and to its children (when connected to them).

Prerequisites

To add a child to a parent in your hierarchy, one of the following two scenarios must be true:

- The management server is installed as a **network service:** Before a parent/child link can be established without the acceptance from the administrator of the child, both computers involved (parent and child) must be added as a user to the other's system administrator role. Refer to Assign and remove users and groups to/from roles.
- The management server is installed as a **user account:** This user account must be a member of the administrator group of the server being linked to before a parent/child link can be established without the acceptance from the administrator of the child.

If neither of these criteria are met, the administrator of the child needs to accept the request for inclusion in the hierarchy (see "Accept inclusion in hierarchy" on page 221) before the link can be established. Refer to OnSSI Federated Architecture overview (see "About OnSSI Federated Architecture" on page 212) for more details.

Add site to hierarchy

1. In the *Management Client* window, in the Federated Sites Hierarchy pane (see "Panes Overview" on page 33), select the relevant site, right-click, and click *Add Site to Hierarchy*.

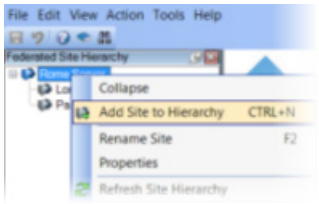

2. Insert the URL of the requested child in the *Add Site to Hierarchy* window.
3. Click *OK*.
4. A link to the new child site is added to the Federated Sites Hierarchy pane.
5. If you can establish the new child link without requesting acceptance from the administrator (see *Prerequisites* described earlier), skip to step 7.

   If **not**, the new child has the awaiting acceptance icon see "Accept inclusion in hierarchy" on page 221and its administrator must authorize the request.
6. Make sure the child's administrator authorizes the link request (this is done from the child site).
7. The new parent/child link is established and the Federated Sites Hierarchy pane is updated with the icon for the new child.

> *Due to synchronization issues, any changes made to children located far from your home-site might take some time to be reflected in the Federated Sites Hierarchy pane. Refer to Basic rules of federated sites (on page 214).*

Accept inclusion in hierarchy

You must accept a child link request manually if your administrator settings require this.

- If the management server is installed as a **network service:** Computers involved must **not** be added as users to each other's Ocularis system administrator role, but should be added as another **non**-administrator role. Refer to Assign and remove users and groups to/from roles.
- If the management server is installed as a **user account:** This user account must **not** be a member of the administrator role of the server being linked to.

Otherwise inclusion will take place automatically.

Also refer to Administrator role and federated sites (see "Administrators role and federated sites" on page 215).

Prerequisites

The potential child must have received a link request from the potential parent. As a result, the child has the awaiting acceptance icon.

Accept inclusion in hierarchy

1. In the Management Client window (of the potential child), in the Federated Sites Hierarchy pane (see "Panes Overview" on page 33), select the relevant site, right-click, and click ***Accept Inclusion in Hierarchy***.

2. Click *Yes*.
3. The new parent/child link is established and the Federated Sites Hierarchy pane is updated with the normal site icon for the selected site.

> *Due to synchronization issues, any changes made to children located far from your home-site might take some time to be reflected in the Federated Sites Hierarchy pane. Refer to Basic rules of federated sites (on page 214),* Synchronization of Hierarchy.

Connect to another site in hierarchy

You can connect to all sites in the federated architecture if your administrator settings are correct.

Prerequisites

To connect from one site in your hierarchy to another, one of the following two scenarios must be true:

- The management server is installed as a **network service:** Both computers involved must be added as users to each other's Ocularis system administrator role. Refer to Assign and remove users and groups to/from roles.
- The management server is installed as a **user account:** This user account must be a member of the administrator group of the server being linked to.

Refer to Administrator role and federated sites section (see "Administrators role and federated sites" on page 215).

Connect to another site in hierarchy

Click the wanted site in the Federated Site Hierarchy pane (see "Panes Overview" on page 33). A brief dialog informs you that you are being connected to the new site. When connection is complete, your view in the Federated Sites Hierarchy pane will change to reflect that you are connected to a different site.

In this example, the user was logged into the home-site *Rome Server* and next connects to the child *Paris Server*:

Do I log out of my home-site when I connect to another site in the hierarchy? Both yes and no. Embedded in your home-site log-in is an automated and seamless log-in to its children as well, using the same credentials as your home-site log-in. However, even though you technically log out of your home-site when connecting to one of its children, you still see the world as your (former) home-site sees it.

Detach a site from hierarchy

Detaching/Removing a site from its hierarchy involves two different results **depending on where in the federated architecture you are located**.

If you are within your hierarchy-except your home-site-this will detach the selected site from the rest of the hierarchy. You will no longer be able to see the detached site.

If, on the other hand, you are located at your home-site, your home site will be detached from the rest of the hierarchy including any sites located under your home-site. Your home-site becomes the new top-site.

Detach child from hierarchy **(Location: Any site)**

Prerequisites

The site you are detaching is any site, **except** your home-site.

To detach child from a hierarchy

1. In the Management Client window, in the Federated Sites Hierarchy pane (see "Panes Overview" on page 33), right-click the site you want to detach-**except** the home-site-select *Detach Site from Hierarchy*.

2. Click *Yes*.
3. The detached site is removed and the *Federated Sites Hierarchy* pane is updated.

Tip : You do not have to connect to a site to detach it. Just point your mouse to the relevant site and right click, select *Detach Site from Hierarchy*.

Detach home-site from parent hierarchy **(Location: Home-site, which has a parent)**

Prerequisites

Your home-site must be the child of another site, i.e. have a parent.

To detach home-site from a parent hierarchy

1. In the *Management Client* window, in the Federated Sites Hierarchy pane (see "Panes Overview" on page 33), right-click the **home-site**, and click *Detach Site from Hierarchy*.
2. Click *Yes*.
3. The *Federated Sites Hierarchy* pane is updated, your home-site becomes the new top-site, and the normal site icon  changes to a top-site  icon.
4. Click *OK*.

Due to synchronization issues, changes might take a little time to be reflected in the Federated Sites Hierarchy pane (see "Panes Overview" on page 33). Refer to Basic rules of federated sites *(on page 214)*.

Refresh site hierarchy

Automatic synchronizations happen regularly through all steps of your parent/child setup. But if you want a current overview of things, and do not want to wait for the next automatic synchronization, you can refresh.

When refreshing, the home-site will display a current overview of the state of things from the home-site's point-of-view.

Note that only changes saved by the home-site since the last synchronization will be reflected—changes further down in the hierarchy will not be reflected. For this, a full scheduled synchronization is needed.

1. In the Management Client window, in the Federated Sites Hierarchy pane (see "Panes Overview" on page 33), right-click the home-site, and click select *Refresh Site Hierarchy*.
2. The Federated Sites Hierarchy pane is refreshed, reflecting any changes.

It is not possible to schedule your own synchronizations.

Rename site

You can rename both your home-site and its children when connected to them.

1. In the Management Client, in the Federated Sites Hierarchy pane (see "Panes Overview" on page 33), select the relevant site, right-click, and click *Rename Site*.
2. You can now overwrite the name of the site.
3. The Federated Sites Hierarchy pane is updated, reflecting the name-change.

Due to synchronization issues, any changes to remote children might take some time to be reflected in the Federated Sites Hierarchy pane. Refer to Basic rules of federated sites (on page 214).

Set site properties

You can view and, possibly, edit properties on your home-site and its children.

1. In the Management Client, in the Federated Sites Hierarchy pane (see "Panes Overview" on page 33), select the relevant site, right-click, and select *Properties.*

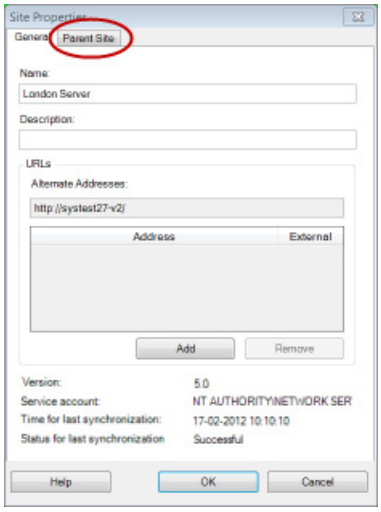

2. If needed, change the following:

General tab

Information related to the site you are currently connected to:

- ***Name****:* Enter the name of the site displayed in the Federated Sites Hierarchy pane and the Site Navigation pane (see "Panes Overview" on page 33).
- ***Description****:* Enter a description of the site.
- ***URLs****:* Use the list to add and remove URL(s) for this site and indicate whether they are external or not.
- ***Version:*** Version number of the site/management server.
- ***Service account:*** The service account under which the management server is running.
- ***Time for last synchronization:*** Last synchronization date.
- ***Status for last synchronization:*** Status of last synchronization. It can be either *Successful* or *Failed*. If failed, further information is offered.

Click *OK* to save changes.

Parent Site tab **(available on child sites only—marked in red)**

Non-editable information regarding the parent of the child you are currently connected to:

- ***Name****:* Shows the name of the parent to be displayed in the Federated Sites Hierarchy pane and Site Navigation pane (see "Panes Overview" on page 33).
- ***Description****:* Shows a description of the parent.
- ***URLs****:* Lists URL(s) for this parent and indicates whether they are external or not.
- ***Version:*** Version number of the site/management server.
- ***Service account:*** The service account under which the management server is running.
- ***Time for last synchronization:*** Last synchronization date.
- ***Status for last synchronization:*** Status of last synchronization. It can be either *Successful* or *Failed*. If failed, further information is offered.

Due to synchronization issues, any changes made to remote children might take some time to be reflected in the Site Navigation pane (see "Panes Overview" on page 33). Refer to Basic rules of federated sites (on page 214).

Backup, restore and move system configuration

Scheduled backup and restore of system configuration

Regularly backing up your system database is always recommended—especially if you have a larger system setup. Having a scheduled regular backup provides you with an always up to date backup. In case of a disaster recovery scenario, regular backups limit your loss of data to what was changed since last backup. Furthermore, it offers you the ability to quickly restore your system configuration. However, regularly backing up also has the added benefit that it flushes your Microsoft® SQL Server's transaction log.

If you have a smaller setup and do not feel the need for regular scheduled backup, refer to Manual backup and restore of system configuration (on page 228).

The management server stores your system's configuration in a database. When backing up/restoring management server(s), make sure that this database is included in the backup/restore.

Flush SQL server transaction log

What is the SQL server transaction log and why does it need to be flushed? Each time a change in the system's data occurs, the SQL Server will log this change in its transaction log - regardless whether it is a SQL Server on your network or a SQL Server Express edition. The transaction log is essentially a security feature that makes it possible to roll back and undo changes to the SQL Server database. The SQL Server by default stores its transaction log indefinitely, and therefore the transaction log will over time build up more and more entries. The SQL Server's transaction log is by default located on the system drive, and if the transaction log just grows and grows, it may in the end prevent Windows from running properly. Flushing the SQL Serve's transaction log from time to time is a good idea. However, flushing it does not in itself make the transaction log file smaller, but it prevents it from growing out of control. Your system does not, however, automatically flush the SQL Server's transaction log at specific intervals. You can also do several things on the SQL Server itself to keep the size of the transaction log down. For numerous articles on this topic, go to support.microsoft.com and search for SQL Server transaction log.

Prerequisites

SQL Server Express Edition users only: Microsoft® SQL Server Management Studio Express, a tool download-able for free from www.microsoft.com/downloads.

Among its many features for managing SQL Server Express databases are some easy-to-use backup and restoration features. Download and install the tool on your management server.

Scheduled back up of system configuration

1. From Windows' *Start* menu, open Microsoft® SQL Server Management Studio Express by selecting *All Programs > Microsoft SQL Server 2008 > SQL Server Management Studio Express*.
2. In the tool do the following:
   When connecting, specify the name of the required SQL Server. Use the account under which the database was created.
   - Find the *Surveillance* database, containing your entire system configuration, including recording servers, cameras, inputs, outputs, users, rules, patrolling profiles, etc.

     No Surveillance database? Surveillance is the default name of the database containing the system configuration. If you can find the database, but it is not called Surveillance, it could be because you gave the database another name during the management server installation.

     View example... We will assume that the database uses the default name.

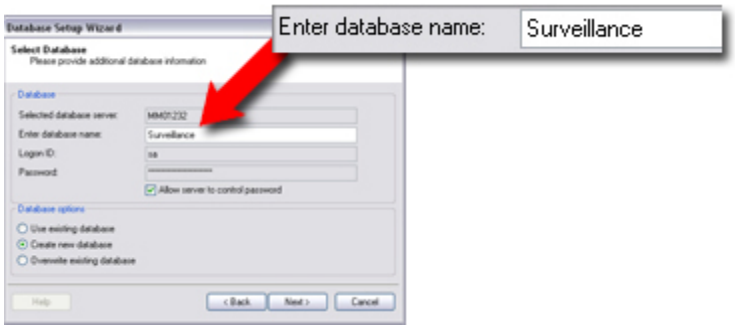

Example: During management server installation it is possible to change the database name from the default name Surveillance to another name

- Make a backup of the *Surveillance* database and make sure to:
  - Verify that the selected database is *Surveillance*
  - Verify that the backup type is **full**
  - Set the schedule for the recurrent backup
  - Verify that the suggested path is satisfactory or select alternative path
  - Select to **verify backup when finished** and to **perform checksum before writing to media**.

3. Follow the instructions in the tool to the end.

Tip: Also consider backing up the *SurveillanceLog* database, using the same method.

Back up log server database

Handle the *SurveillanceLogServer* database using the same method as when handling system configuration described earlier in this topic. The *SurveillanceLogServer* database (name may be different if you renamed the system configuration database) contains all your system logs, including errors reported by recording servers and cameras.

The database is located where the Log Server Service is installed, typically the same place as your management server. Backing up this database is not vital since it does not contain any system configuration, but you may later appreciate having access to system logs from before the management server backup/restore.

Restore system configuration (from scheduled back up)

Prerequisite: To prevent configurational changes being made while you restore the system configuration database, stop the:

- Management Server service (see "Management Server service and Recording Server service" on page 244)
- Event Server Service (can be done from Windows *Services* (search for *services.msc* on your machine. Within *Services*, locate *OnSSI Ocularis Event Server*))
- World Wide Web Publishing Service, also known as the Internet Information Service (IIS). Learn how to stop the IIS at: http://technet.microsoft.com/en-us/library/cc732317(WS.10).aspx.

Open Microsoft® SQL Server Management Studio Express from Windows' *Start* menu by selecting *All Programs > Microsoft SQL Server 2008 > SQL Server Management Studio Express*.

1. In the tool do the following:
   - When connecting, specify the name of the required SQL Server. Use the account under which the database was created.

- Find the *Surveillance* database, containing your entire system configuration, including recording servers, cameras, inputs, outputs, users, rules, patrolling profiles, etc.
- Make a restore of the *Surveillance* database and make sure to:
  - Select to backup **from** device
  - Select backup media type **file**
  - Find and select your backup file *Surveillance.bak*
  - Select to **overwrite the existing database**.

2. Follow the instructions in the tool to the end.

If you also backed up the *SurveillanceLog* database from the old management server, restore it on the new management server using the same method.

Note that the system basically does not work while the Management Server service (see "Management Server service and Recording Server service" on page 244) is stopped; it is important to remember to start the services again once you have finished restoring the database.

Manual backup and restore of system configuration

Backing up your system database is always recommended. In case of a disaster recovery scenario this offers you the ability to quickly restore your system configuration. Furthermore, being able to easily do a manual backup of your entire system configuration via your Management Client (no need for third-party tools) offers you flexibility, security and full control of your configuration.

The type of backup described in this topic is best suited if you have a smaller system setup and wish to do a one-time, non-scheduled backup. Besides manual backups, it is strongly recommended to also configure regular, scheduled system backups (see "Scheduled backup and restore of system configuration" on page 226)—especially if you run a larger system setup.

Before backing up and restoring any system configuration, you must set a backup folder for this purpose.

1. Right-click the notification area's management server service icon and select *Select shared backup folder...*
2. In the window that appears, browse to the wanted file location.
3. Click *OK* twice.
4. If asked if you want to delete files in the current backup folder, click *Yes* or *No* depending on your needs

Important information:

- Both the user installing and the user doing the restore must be local administrator on the management server **and** on the SQL server.
- Except for your recording servers, your system will be completely shut down for the duration of the restore, which might take some time.
- A backup can only be restored on the system installation where it was created. Furthermore, make sure that the setup is as similar as possible to when the backup was made. Otherwise, the restore might fail.
- If restoring fails during the validation phase, it **will** be possible to start the old configuration again (since no change have been committed).
  If restoring fails elsewhere in the process, rolling back to the old configuration is impossible.
  As long as the backup file is not corrupted, it **will** however be possible to do another restore.
- Restoring replaces the current configuration. This means that any configurational changes since last backup is lost.
- No logs (including audit logs (see "Manage logs" on page 197)) are restored
- Once restoring has started, it cannot be canceled.

Restoring:

1. Right-click the notification area's Management Server service icon and select *Restore Configuration....*
2. Next, you are presented with an important note. Read the contents of the note. Click *Restore.*
3. In the file open dialog, browse to the location of the configuration backup file, select it, and click *Open.*
4. The *Restore Configuration* window will now run, showing progress and status information. Wait for it to finish and click *Close*. Your restore is finished.

Select shared backup folder

Before backing up and restoring any system configuration, you must set a backup folder for this purpose.

1. Right-click the notification area's management server service icon and select *Select shared backup folder...*
2. In the window that appears, browse to the wanted file location.
3. Click *OK* twice.
4. If asked if you want to delete files in the current backup folder, click *Yes* or *No* depending on your needs

Manual back up of system configuration

Important information:

- Your system stays online.
- A backup cannot be used for copying configurations (see "Move system configuration to new management server" on page 230) to other systems.
- Depending on your system configuration, your hardware, and on whether your SQL server, management server and Management Client are installed on the same machine or not, backing up configuration might take some time.
- Logs (including audit logs (see "Manage logs" on page 197)) are **not** part of the configuration backup.

Back up:

All relevant system configuration files will be combined into one single .cnf file, which is saved at a specified location.

1. From the Management Client's menu bar, select *File, Backup Configuration....*
2. Next, you are presented with an important note. Read the contents of the note. Click *Backup.*
3. In the file save dialog, browse to the location where you want to store the configuration backup. Specify a suitable file name, and click *Save.*
4. Let the *Backup Configuration* window finish. Click *Close.* Your backup is finished.

Restore system configuration (from manual back up)

Important information:

- Both the user installing and the user doing the restore must be local administrator on the management server **and** on the SQL server.
- Except for your recording servers, your system will be completely shut down for the duration of the restore, which might take some time.
- A backup can only be restored on the system installation where it was created. Furthermore, make sure that the setup is as similar as possible to when the backup was made. Otherwise, the restore might fail.

- If restoring fails during the validation phase, it **will** be possible to start the old configuration again (since no change have been committed).
  If restoring fails elsewhere in the process, rolling back to the old configuration is impossible.
  As long as the backup file is not corrupted, it **will** however be possible to do another restore.
- Restoring replaces the current configuration. This means that any configurational changes since last backup is lost.
- No logs (including audit logs (see "Manage logs" on page 197)) are restored
- Once restoring has started, it cannot be canceled.

Restoring:

1. Right-click the notification area's Management Server service icon and select *Restore Configuration....*
2. Next, you are presented with an important note. Read the contents of the note. Click *Restore.*
3. In the file open dialog, browse to the location of the configuration backup file, select it, and click *Open.*
4. The *Restore Configuration* window will now run, showing progress and status information. Wait for it to finish and click *Close*. Your restore is finished.

Move system configuration to new management server

It can sometimes be necessary to move the management server installation from one physical server to another. The management server stores your system configuration in a database. If you are moving the management server from one physical server to another, it is vital that you make sure that your new management server also gets access to this database. The system configuration database can be stored in two different ways:

- **Network SQL Server:** If you are storing your system configuration in a database on an existing SQL 2005 or 2008 Server on your network, you can point to the database's location on that SQL Server when installing the management server software on your new management server. In that case, only the following paragraph about management server hostname and IP address applies and you should ignore the rest of this topic:

  Management server hostname and IP address: When you move the management server from one physical server to another physical server, it is by far the easiest to give the new server the same hostname and IP address as the old one. This is due to the fact that the recording server will connect to the hostname and IP address of the old management server. In case the new management server has been given a new hostname and/or IP address, the recording server will not be able to find the management server. Manually stop each recording server in your system, change their management server URL, and when done, restart them.
- **SQL Server Express Edition:** If you are storing your system configuration in a SQL Server Express Edition database on the management server itself, it is important that you back up the existing management server's system configuration database before the move. By backing up the database, and subsequently restoring it on the new server, you will avoid have to reconfigure your cameras, rules, time profiles, etc. after the move.

Some of this prerequisite information is only relevant for users of SQL Server Express Edition. **If you use any other SQL setup, ask your IT department for backup details.**

Prerequisites

- **Your software installation file for installation on the new management server**.
- **Your initial license (.lic) file**, i.e. the one you used when initially installing your system, not the .lic file which is the result of your license activation (see "Activate (Register) Licenses - Online or Offline" on page 44). License activation is, among other things, based on the specific hardware on which the activation took place; therefore an activated .lic file cannot be reused when moving to a new server. Note that if you are also upgrading your system software in connection with the move, you will have received a new initial .lic file together with your new Software License Code (SLC).
- **SQL Server Express Edition users only: Microsoft® SQL Server Management Studio Express**, a tool downloadable for free from www.microsoft.com/downloads. Among its many features for managing SQL Server

Express databases are some easy-to-use backup and restoration features. Download and install the tool on your existing management server *and* on the server which will be your future management server (you will need it for the entire copy process (backup as well as restoration)).

Management server hostname and IP address: When you move the management server from one physical server to another physical server, it is by far the easiest to give the new server the same hostname and IP address as the old one. This is due to the fact that the recording server will connect to the hostname and IP address of the old management server. In case the new management server has been given a new hostname and/or IP address, the recording server will not be able to find the management server. Manually stop each recording server in your system, change their management server URL, and when done, restart them.

Move system configuration:

Moving your system configuration is in reality a three step process:

1. First you make a copy of your system configuration (identical to making a scheduled backup (see "Scheduled back up of system configuration" on page 226))
2. Then you install the new management server on the new server (refer to scheduled backup (see "Scheduled back up of system configuration" on page 226), step 2)
3. And finally you copy/restore your system configuration to the new system (refer to restore a scheduled backup) (see "Restore system configuration (from scheduled back up)" on page 227)

Copy system configuration from old server (step 1)

Prerequisite: Stop the Management Server service (see "Management Server service and Recording Server service" on page 244) to prevent configuration changes being made. This is important since any changes made to the system configuration, between the time you create a copy and the time you restore it on your new management server, will be lost. If changes are made after the copy was made, you will have to make a new copy.

Note that the system basically will not work while the Management Server service (see "Management Server service and Recording Server service" on page 244) is stopped. Remember to start the service again once you have finished backing up the database.

First part of a copy is in reality identical to a scheduled backup (see "Scheduled back up of system configuration" on page 226), steps 1-3.

What happens while the management server is unavailable?

- **Recording servers will still be able to record:** Any currently working recording servers will have received a copy of their configuration from the management server, so they will be able to work and store recordings on their own while the management server is down. Scheduled and motion-triggered recording will therefore work, and event-triggered recording will also work unless based on events related to the management server or any other recording server since these go through the management server.
- **Recording servers will temporarily store log data locally:** They will automatically send log data to the management server when the it becomes available again.
    - **Clients will not be able to log in:** Client access is authorized through the management server. Without the management server, clients will not be able to log in.
    - **Already logged in clients can remain logged in for up to an hour:** When clients log in, they are authorized by the management server and can communicate with recording servers for up to one hour. If you can get the new management server up and running within an hour, many of your users will not be affected.
    - **No ability to configure the system:** Without the management server, you will not be able to change system configuration.

Even though some users might not experience loss of contact, we recommend that you inform your users about the risk of losing contact with the surveillance system while the management server is down.

Copy log server database

Handle the *SurveillanceLogServer* database using the same method as when handling system configuration described earlier in this topic. The *SurveillanceLogServer* database (name may be different if you renamed the system configuration database) contains all your system logs, including errors reported by recording servers and cameras.

The database is located where the Log Server Service is installed, typically the same place as your management server. Backing up this database is not vital since it does not contain any system configuration, but you may later appreciate having access to system logs from before the management server backup/restore.

Install new management server on new server (step 2)

Installing a management server is divided into three steps. During step 2 of the installation on your new management server, make sure you select *Create a new database* for the system configuration database, even though you have a backup of the database from your old management server.

Next (see "Copy/restore system configuration to new server (step 3)" on page 232), overwrite the new and empty database by restoring the backup we just created. Since you are going to overwrite the new and empty database, it is important that it has the same name as the backed-up database (if your backed-up database has the default name Surveillance, just use the default name *Surveillance* when creating the new database too).

The password for the database is not significant in this backup/restore context, but we recommend that you just use the default setting *Allow server to control password*.

Copy/restore system configuration to new server (step 3)

Prerequisite: To prevent configurational changes being made while you restore the system configuration database, stop the:

- Management Server service (see "Management Server service and Recording Server service" on page 244)
- Event Server Service (can be done from Windows *Services* (search for *services.msc* on your machine. Within *Services*, locate *OnSSI Ocularis Event Server*))

World Wide Web Publishing Service, also known as the Internet Information Service (IIS). Learn how to stop the IIS at: http://technet.microsoft.com/en-us/library/cc732317(WS.10).aspx.

This should all be done to the **new** management server.

Since second part of a copy is in reality identical to a restore, refer to Restoring system configuration (from scheduled back up) (see "Restore system configuration (from manual back up)" on page 229), steps 1-2 and rest of the topic for details.

Device drivers

Manage video device drivers

Video device drivers are modules used for controlling/communicating with the camera devices connected to a recording server. The video device drivers should therefore be installed on each recording server on your system. Video device drivers are installed automatically during the initial installation of your system. However, new versions of video device drivers are released and made available on our website: www.onssi.com.

Remove video device drivers

Video device drivers are modules used for controlling/communicating with the camera devices connected to a recording server. When the video device drivers are removed, communication between the recording server and the camera devices will no longer be possible.

To remove video device drivers use the following procedure on the recording server computer on which the video device drivers are installed:

The following procedure describes standard system component removal in recent Windows versions; the procedure may be slightly different in older Windows versions:

1. In Windows' *Start* menu, select *Control Panel*, and then...
   - If using *Category* view, find the *Programs* category, and click *Uninstall* a program.
   - If using *Small icons* or *Large icons* view, select *Programs and Features*.
2. In the list of currently installed programs, right-click the required OnSSI program or service.
3. Select *Uninstall* if you wish to uninstall all components. Select *Change* if you only wish to uninstall some components
4. Follow the removal instructions.

Note that you should not remove the device pack when upgrading, you can install the new version on top of the old one. The device pack should only be removed when whole system is uninstalled.

Failover recording servers—regular/hot standby

About failover recording servers—regular and hot standby

Depending on the recording component, functionality described here may be limited or unavailable.

A failover recording server is a spare recording server which takes over from a normal recording server in case this becomes unavailable. In the following, the term *failover recording server* is used as an umbrella term for both *regular failover recording servers* and *hot standby servers*.

You can configure a failover recording server in two ways, as a **regular failover recording server** or as a **hot standby server** (see "**Assign failover recording servers**" on page 74). In a regular failover setup, a failover recording server can be grouped with other failover recording servers in a failover group. The entire failover group is dedicated to taking over from any of several preselected recording servers, should one of these become unavailable. A failover recording server in a **hot standby setup** is dedicated to take over from **one** recording server only. Because of this, they can be kept in a "standby" mode which means that they are already started with the correct/current configuration of "their" recording server and are ready to take over more quickly than a regular failover recording server. As mentioned, hot standby servers are assigned to one recording server only and therefore cannot be grouped. Likewise, regular failover servers already part of a failover group cannot be selected as hot standby servers.

A failover group can contain one or more regular failover recording servers. Grouping (see "Group failover recording servers" on page 238) has a clear benefit: when you later specify which failover recording servers should be able to take over from a recording server, you select a group of failover recording servers. If the selected group contains more than one failover recording server, this offers you the security of having more than one failover recording server ready to take over if a recording server becomes unavailable. You can create as many failover groups as required and group them as needed. However, a failover recording server can only be a member of one group at a time.
Failover recording servers in a failover group are ordered in sequence. This sequence determines in which order the failover recording servers should take over from a recording server when needed. By default, this sequence will reflects the order in which the failover recording servers have been incorporated in the failover group—first in, is first in sequence—but this can easily be changed.

Failover recording servers are installed like regular recording servers; refer to Install failover recording server (see "Install failover recording server (recording server)" on page 16). Once failover recording servers are installed, they automatically become visible in the Management Client.

Tip: If a new failover recording server does not become visible in the Management Client, verify that the failover recording server has been configured with the correct IP address/hostname of the management server. Also verify that the user account under which the Failover Server service runs has access to your system with administrator rights.

By default failover recording servers are enabled. If you have disabled it, you must enable it (see "Setup and enable failover recording servers" on page 237) before it can take over from recording servers.

All failover recording servers should always be installed on separate computers.

It is possible to specify what type of failover support you want on device-level (see "Assign failover recording servers" on page 74). For each device on a recording server you can select full, live only or no failover support. This helps you prioritize your failover resources and, for example, only set up failover for video and not for audio, or only have failover on essential cameras, not on less important ones.

A failover recording server has two services installed:

- A Failover Server service, which handles the processes of taking over from the recording server. By default, this service is always running, constantly checking the state of relevant recording servers.
- A Failover Recording Server service (on page 240), which enables the failover recording server to act as a recording server.
  In a failover group setup, this service is only started when required, i.e. when the regular failover recording server should take over from the recording server. Starting this service typically takes a couple of seconds, but may take longer depending on local security settings, etc.
  In a hot standby setup, this service is always running, allowing the hot standby server to take over faster than the regular failover recording server.

Illustration: Failover process in details

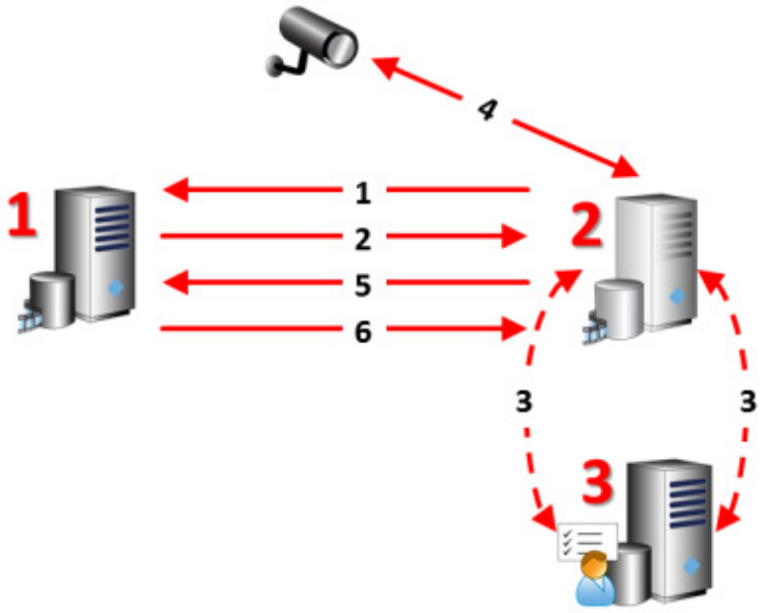

Involved **servers** (numbers in red):

1. Recording server
2. Failover recording server
3. Management server.

Regular failover setups:

1. To check whether it is running or not, a failover recording server has a non-stop TCP connection to a recording server.
2. This connection is interrupted, i.e. the recording server is not running.
3. The failover recording server requests the current configuration of the recording server from the management server. The management server sends the requested configuration, the failover recording server receives the configuration, starts up, and starts recording on behalf of the recording server.
4. The failover recording server and the relevant camera(s) exchange video data.
5. The failover recording server continually tries to re-establish connection to the recording server.
6. When the connection to the recording server is re-established, the failover recording server shuts down and the recording server fetches video data (if any) recorded during its down-time and the video data is merged back in to the recoding servers database.

Failover steps for **hot standby** setups:

1. To check whether it is running or not, a hot standby server has a non-stop TCP connection to its assigned recording server.
2. This connection is interrupted, i.e. the recording server is not running.
3. From the management server, the hot standby server already knows the current configuration of its assigned recording server and starts recording on its behalf.
4. The hot standby server and the relevant camera(s) exchange video data.
5. The hot standby server continually tries to re-establish connection to the recording server.
6. When the connection to the recording server is re-established and the hot standby server goes back to hot standby mode, the recording server fetches video data (if any) recorded during its down-time and the video data is merged

back in to the recoding servers database.

FAQs: failover recording servers

How does a failover recording server know when to take over? It polls (i.e. regularly checks the state of) relevant recording servers every 0.5 seconds. If a recording server does not reply within 5 seconds, the recording server is considered unavailable and the failover recording server takes over.

How long does it take for a failover recording server to take over? For a regular failover server it takes approximately 5 seconds plus the time it takes for the failover recording server's Recording Server service to start.
A hot standby server however, can do it faster since it is already in hot standby mode and only has to start its cameras to deliver feeds.
During the start up period, it will not be possible to store recordings, neither will it be possible to view live video from affected cameras.

What happens when a recording server becomes available again? It will automatically take over from the failover recording server, and recordings stored by the failover recording server will automatically be merged into the standard recording server's databases. How long the merging process takes depends on the amount of recordings, on network capacity, etc. During the merging process, it will not be possible to browse recordings from the period during which the failover recording server took over.

What if a failover recording server must take over from another recording server during the merging process? In a regular failover recording server setup, it will postpone the merging process with recording server A, and take over from recording server B. When recording server B becomes available again, the regular failover recording server will take up the merging process with recording server A, after which it will begin merging with recording server B.
In a hot standby setup, a hot standby server cannot take over for another recording server because it can only be hot standby for a single recording server. But if that recording server fails again, the hot standby will just take over again and also keep the recordings from the previous period. Recordings are kept until they are merged back to the primary recorder or until the failover recording server runs out of disk space.

Will I lose recordings? A failover solution does not provide complete redundancy. It is, however, a very reliable way of minimizing downtime. When the recording server becomes available again, the Failover Server service will make sure that the recording server is ready to store recordings again. Only then is the responsibility for storing recordings handed back to the standard recording server. So, a loss of recordings at this stage of the process is very unlikely.

How will clients experience a failover setup? Clients should hardly notice that a failover recording server is taking over. There will however be a short break—usually only some seconds—when the failover recording server is taking over. During this break there will be no access to video from the affected recording server. Clients will be able to view live video as soon as the failover recording server has taken over. Since recent recordings are stored on the failover recording server, they will also be able to play back recordings from after the failover recording server took over. Clients will not be able to play back older recordings stored only on the affected recording server until that recording server is functioning again and has taken over from the failover recording server. It will be possible to access archived recordings stored at accessible locations (i.e. on available network drives), but not archived recordings stored at inaccessible locations (i.e. on the unavailable recording server itself or on an unavailable network drive). When the recording server is functioning again, there will usually be a merging process during which failover recordings are merged back into the recording server's database. During this process, it will not be possible to play back recordings from the period during which the failover recording server took over.

Is there a failover solution for failovers? In a regular failover setup, setting up one failover recording server as backup for another failover recording server is not necessary. This is because you do not allocate particular failover recording servers to take over from a standard recording server; rather you allocate failover groups. A failover group must contain at least one failover recording server, but you can add as many failover recording servers as needed. Provided a failover group contains more than one failover recording server, there will be more than one failover recording server capable of taking over.
In a hot standby setup, it is also not possible to set up a failover recording servers or hot standby servers for a hot standby server.

For more information about failover setups, refer to the description of the Management Client's Failover tab (see "Failover tab (recording server properties)" on page 74).

Install failover recording servers

For information about installing failover recording servers, refer to Install failover recording server (see "Install failover recording server (recording server)" on page 16).

Setup and enable failover recording servers

If you have disabled the failover recording server, you must enable it before it can take over from standard recording servers.

Do the following to enable a failover recording server and edit its basic properties:

1. In the Site Navigation pane (see "Panes Overview" on page 33), select *Servers*, *Failover Servers*. This opens a list of installed failover recording servers and failover groups.
2. In the Overview pane (see "Panes Overview" on page 33), select the required failover recording server.
3. Right-click and select *Enabled*. The failover recording server is now enabled.

Tip: You can tell a failover recording server's status from its icon. The server in the example above has a green tick mark, indicating is either waiting or "watching". Furthermore, by hovering over it, a tooltip appears containing the text entered in the ***Description*** field of the failover recording server. You may want to use this to indicate which recording server the failover recording server is configured to take over from.

4. To edit failover recording server properties (see "Failover recording server properties" on page 237), go to the *Info* tab:

Info tab of a failover recording server

5. When done, go to the *Network* tab. Here you can define the failover recording server's public IP address, etc. This is relevant especially if using NAT (Network Address Translation) and port forwarding. Refer to the standard recording server's *Network* tab (see "Network tab (recording server properties)" on page 78) for more information.

Failover recording server properties

- ***Name:*** Name as it appears in the Management Client, in logs, etc.
- ***Description:*** Optional description, for example which recording server it is taking over from or a description of the server's physical location.

- ***Host name:*** Non-editable field displaying the network address of the failover recording server.
- ***UDP port:*** The port number used for communication between failover recording servers. By default, port 8844 is used.
- ***Database location:*** Specify the path to the database used by the failover recording server for storing recordings.

> *The database path cannot be changed while the failover recording server is taking over from a recording server. Changes will be applied when the failover recording server is no longer taking over from a recording server.*

- ***Enable this failover server:*** Clear to disable the failover recording server (by default selected). Note that failover recording servers must be disabled to be able to take over from recording servers.

Group failover recording servers

1. In the Site Navigation pane (see "Panes Overview" on page 33), select *Servers*, *Failover Servers*. This opens a list of installed failover recording servers and failover groups.
2. In the Overview pane (see "Panes Overview" on page 33), right-click the top-node *Failover Groups* and select *Add Group* from the menu that appears.
3. Specify a name (in this example *Failover Group 1*) and a description (optional) of your new group. Click *OK*.
4. Right-click the group (*Failover Group 1*) you just created. From the menu that appears, select *Edit Group Members*. This opens the *Select Group Members* window.
5. Drag and drop or use the buttons to move the selected failover recording server(s) from the left to the right side:

Click *OK*. The selected failover recording server(s) now belongs to the group (*Failover Group 1*) you just created:

6. Next, go to the ***Sequence*** tab. Click ***Up*** and ***Down*** to set the internal sequence of the regular failover recordings servers in the group.

Failover group properties

The *Info* tab:

- ***Name:*** Name as it appears in the Management Client, in logs, etc.

- ***Description:*** Optional description, for example a description of the server's physical location.

The *Sequence* tab:

- ***Specify the failover sequence:*** Use ***Up*** and ***Down*** to set the wanted sequence of regular failover recording servers within the group.

Assign failover recording servers

On the **Failover** tab of a recording server, you can choose between 3 different types of failover setups:

a No failover setup

b A primary/secondary failover setup

c A hot standby setup.

If you select **b** and **c**, you must select the specific server/groups. With **b**, you must also select a primary and optionally a secondary failover group. If the recording server becomes unavailable, a failover recording server from the primary failover group will take over. If you have also selected a secondary failover group, a failover recording server from the secondary group will take over in case all failover recording servers in the primary failover group are busy. This way you only risk not having a failover solution in the rare case when all failover recording servers in the primary, as well as in the secondary, failover group are busy.

1. In the Site Navigation pane (see "Panes Overview" on page 33), select *Servers*, *Recording Servers*. This opens a list of recording servers.
2. In the Overview pane (see "Panes Overview" on page 33), select the wanted recording server, go to the ***Failover*** tab.
3. To choose failover setup type (see "About failover recording servers—regular and hot standby" on page 234), select either ***None***, ***Primary failover server group***/***Secondary failover sever group*** or ***Hot standby server***. If relevant, select the needed server or groups from the dropdowns.

   You cannot select the same failover group as both primary and secondary failover group. Also regular failover servers already part of a failover group cannot be selected as hot standby servers.

   Tip: From the ***Primary***/***Secondary failover server group*** dropdowns, select ***Add new...*** to create new failover groups and add failover recording servers.
4. Next, click ***Advanced failover settings...***, this opens the ***Advanced Failover Settings*** window listing all devices attached to the selected recording server.

   Tip: Even if you selected ***None***, ***Advanced failover settings*** will be available. Any selections are kept for later failover setups.
5. To specify the level of failover support, select ***Full Support***, ***Live Only*** or ***Disabled*** for each device in the list. Click ***OK***.
6. Finally, in the ***Failover service communication port (TCP)*** field, edit the port number if needed.

Failover tab properties

- ***None:*** Select a setup without failover.
- ***Primary failover server group / Secondary failover sever group:*** Select a regular failover setup with one primary and possibly one secondary failover server group. Also, from the attached dropdown, select a primary failover group and possibly a secondary failover group.
- ***Hot standby server:*** Select a hot standby setup. Also, from the dropdown, select a hot standby server.
- ***Advanced failover settings...:*** Opens the ***Advanced Failover Settings*** window.
  - ***Full Support:*** Select to get full failover support for the device.

  - ***Live Only:*** Select to get live failover support for the device.
  - ***Disabled:*** Select to disable failover support for the device.

- ***Failover service communication port (TCP):*** By default, the port number is 11000. This port is used for communication between recording servers and failover recording servers. If changed, the recording server in question **must** be running and **must be** connected to the management server meanwhile.

Failover-related events

Your system features two failover-related events, *Failover Started* and *Failover Stopped*, which you can use when creating rules (see "Manage rules" on page 165). The two events are further described in the Events overview (on page 161).

Read failover recording server status icons

The following icons represent the status of failover recording servers (icons are visible in the Management Client, in the Overview pane (see "Panes Overview" on page 33)):

Icon	Description
	The failover recording server is either waiting or "watching". When waiting, the failover recording server is not configured to take over from any recording server yet. When "watching", the failover recording server is configured to watches one or more recording servers.
	The failover recording server has taken over from the designated recording server. **Tip:** When hovering over the server icon in the Management Client, a tooltip appears. You can use this tooltip to identify which recording server the failover recording server has taken over from. The tooltip-text is defined in the failover recording server's Description field (see "Failover recording server properties" on page 237).
	Connection to the failover recording server is broken.

Failover Recording Server service

When you have installed a failover recording server, you are able to check the state of the Failover Recording Server service by looking at the Failover Recording Server service icon in the notification area **of the computer running the failover recording server**. The notification area icon also lets you start and stop the Failover Recording Server service, view status messages, etc.

Tip: The notification area is occasionally also known as the **system tray**, it is located at the far right of the management server computer's Windows taskbar.

Example: Failover Recording Server service icon in notification area; note that failover recording servers also have a Recording Server service (other icon)

While the Failover Recording Server service is stopped, the failover recording server will not be able to take over from standard recording servers.

Start and stop the Failover Recording Server service

The Failover Recording Server service starts automatically. If you have stopped the service manually, you can start and stop it the following way:

1. Right-click the notification area's failover recording server icon.
2. From the menu that appears, select *Start Failover Recording Server service* or *Stop Failover Recording Server service*, depending on your needs.

Change the management server address

The failover recording server must be able to communicate with your system's management server. You therefore specify the IP address/hostname of the management server during the installation of the failover recording server.

Should you later need to change the address of the management server, you do it the following way:

In order to be able to change the management server address, the Failover Recording Server service must be stopped.

1. Stop the Failover Recording Server service (see "Start and stop the Failover Recording Server service" on page 241).
2. Right-click the notification area's Failover Recording Server service icon again.
3. From the menu that appears, select *Change Settings...* The *Failover Recording Server Settings* window appears. You are able to change the following setting:
   - **Management server hostname / IP address:** Lets you specify the IP address (example: 123.123.123.123) or host name (example: ourserver) of the management server with which the failover recording server should be able to communicate.

View status messages

1. Right-click the notification area's *OnSSI Failover Server service* icon.
2. From the menu that appears, select *Show Status Messages.* The *Failover Server Status Messages* window appears, listing time-stamped status messages.

View version information

Knowing the exact version of your *Failover Recording Server service* is an advantage if you need to contact product support.

1. Right-click the notification area's *OnSSI Failover Recording Server service* icon.
2. From the menu that appears, select *About...*
3. A small dialog opens. The dialog will show the exact version of your *Failover Recording Server service.*

Database corruption

Protect recording databases from corruption

If a recording server's databases become corrupted, the recording server is in many cases able to repair the corrupt databases. While the ability to repair corrupt databases is highly valuable, it is of course even better to take steps to ensure that your databases do not become corrupted:

Power outages: Use a UPS

The single-most common reason for corrupt databases is the recording server being shut down abruptly, without files being saved and without the operating system being closed down properly. This may happen due to power outages, due to somebody accidentally pulling out the server's power cable, or similar.

The best way of protecting your recording servers from being shut down abruptly is to equip each of your recording servers with a UPS (Uninterruptible Power Supply).

The UPS works as a battery-driven secondary power source, providing the necessary power for saving open files and safely powering down your system in the event of power irregularities. UPSs vary in sophistication, but many UPSs include software for automatically saving open files, for alerting system administrators, etc.

Selecting the right type of UPS for your organization's environment is an individual process. When you assess your needs, however, bear in mind the amount of runtime you require the UPS to be able to provide if the power fails. Saving open files and shutting down an operating system properly may take several minutes.

Windows Task Manager: Careful when ending processes

When working in Windows Task Manager, be careful not to end any processes which affect the surveillance system. If you end an application or system service by clicking *End Process* in the Windows Task Manager, the process in question will not be given the chance to save its state or data before it is terminated. This may in turn lead to corrupt camera databases.

Windows Task Manager will typically display a warning if you attempt to end a process. Unless you are absolutely sure that ending the process will not affect the surveillance system, make sure you click *No* when the warning message asks you if you really want to terminate the process.

Hard disk failure: Protect your drives

Hard disk drives are mechanical devices, and as such they are vulnerable to external factors. The following are examples of external factors which may damage hard disk drives and lead to corrupt camera databases:

- Vibration (make sure the surveillance system server and its surroundings are stable)
- Strong heat (make sure the server has adequate ventilation)
- Strong magnetic fields (avoid)
- Power outages (make sure you use a UPS (see "Power outages: Use a UPS" on page 242))
- Static electricity (make sure you ground yourself if you are going to handle a hard disk drive).
- Fire, water, etc. (avoid)

SQL database administration

Update SQL server address

When a system is installed as a trial, or if a large installation is restructured, a need for using a different SQL database may arise. This can be handled with the ***Update SQL Server Address*** tool. With it, you can change the addresses of the SQL servers used by the management server and the log server. Only limitation is that you cannot change the management server SQL address at the same time as the log sever SQL address. It is however possible to do so sequentially.

IMPORTANT: This SQL update must be done locally on the machine where the management server **or** log server are installed; you **cannot** do it from the Management Client. If your management server is not located on the same machine, you can still use the tool, but you will have to run it on both the machine where the management server is installed.

Updating management server SQL address

1. If your management server is located:

a) together on the same machine and you wish to update both SQL addresses, go to the machine where your management server is installed.

b) on different machines and you wish to update the management server SQL address, go to the machine where your management server is installed.

2. If:
   - **a** and **b**, go to the notification area of the taskbar (a.k.a. *Systray*). Right-click the ***Management Server*** icon, select ***Update SQL address...***.

3. The ***Update SQL Server Address*** dialog appears. Select ***Management Server*** and click ***Next***.
4. Enter or select the new SQL server and click ***Next***.
5. Select the new SQL database and click ***Select***.
6. Wait while the address change takes place. When a confirmation message is presented, click ***OK***.

Updating log server SQL address

1. If your management server and log server are located:

a) together on the same machine, go to the machine where your management server is installed.

b) on different machines, go to the machine where your management server is installed and copy the directory *%ProgramFiles%\OnSSI\Management Server\Tools\ChangeSqlAddress* (with content) to a temporary directory on another server.

2. If:
   - **a,** go to the notification area of the taskbar (a.k.a. *Systray*). Right-click the ***Management Server*** icon, select ***Update SQL address...***.
   - **b**, paste the directory you copied to a temporary place on the machine where the log server is installed and run the included file:*VideoOS.Server.ChangeSqlAddress.exe*.
3. The ***Update SQL Server Address*** dialog appears. Select ***Log Server*** and click ***Next***.
4. See steps **4**, **5**, and **6** above.

Services administration

Management Server service and Recording Server service

When the management server software is installed, you are able to check the state of the Management Server service by looking at the *Management Server service* icon in the notification area **of the computer running the management server**.

Likewise, when the recording server software is installed, you are able to check the state of the Recording Server service by looking at the *Recording Server service* icon in the notification area of the computer running the recording server in question.

The notification area icon also lets you start and stop the Management Server service/Recording Server service, view status messages, etc.

Tip: The notification area is also known as the *system tray*. It is located at the far right of the management / recording server's Windows taskbar.

IMPORTANT: When the **Recording Server service** is running, it is **very** important that neither Windows Explorer nor other programs are accessing Media Database files or folders associated with your system setup. Otherwise, the recording server might not be able to rename or move relevant media files. Unfortunately, this might bring the recording server to a halt. If this situation has already occurred, stop the Recording Server service, close the program accessing the media file(s) or folder(s) in question, and simply restart the Recording Server service.

Example: *Management Server service* and *Recording Server service* icons in notification area

Access the server service (on page 244)

Start the server service (on page 244)

Stop the server service (on page 245)

Change recording server settings (on page 245)

View status messages (on page 245)

View version information (on page 245)

Work with recording server settings in details (see "Recording server settings" on page 246)

Read server service state icons (see "Read server service icons - management, recording and failover" on page 246)

Access the server service

1. Right-click the notification area's server service icon.
2. From the menu that appears, depending on server type, select the needed icon.

If using multiple instances (see "Multiple recording server instances" on page 25) of the Recording Server service, a sub-menu lets you select whether you want to start a particular instance or all instances.

Start the server service

1. Access the server service (on page 244).
2. Select either *Start Management Server service* or *Start Recording Server service*.

Stop the server service

While the recording server service is stopped, your system will not be able to interact with devices connected to the recording server. Consequently, no live viewing or recording will be possible.

While the management server service is stopped, you will not be able to use the Management Client at all.

1. Access the server service (on page 244).
2. Select either *Stop Recording Server service* or *Stop Management Server service*.

Change recording server settings

To change basic settings for the Recording Server service, such as which port numbers to use, do the following:

To be able to change settings, the Recording Server service must be stopped. While the Recording Server service is stopped, the system will not be able to interact with devices connected to the recording server. Consequently, no live viewing or recording will be possible.

1. Refer to Access the server service (on page 244).
2. Select *Stop Recording Server service*.
3. Right-click the notification area's recording server icon.
4. From the menu that appears, select *Change Settings...*

   The *Recording Server Settings* window (see "Recording server settings" on page 246) appears. Change the appropriate settings.

View status messages

1. Refer to Access the server service (on page 244).
2. Select *Show Status Messages*.

   Depending on the current server type, either the *Management Server Status Messages* or *Recording Server Status Messages* window appears, listing time-stamped status messages:

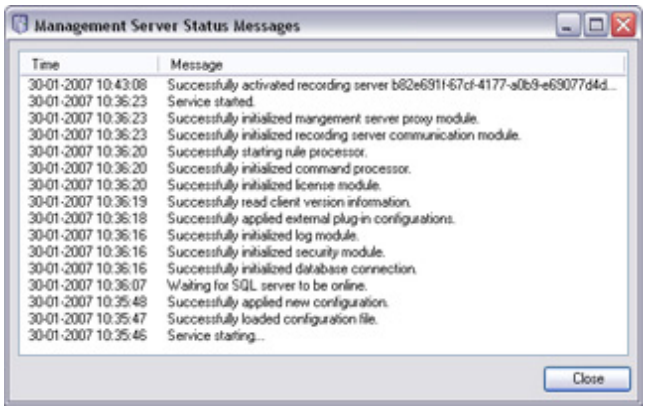

Example from Management Server service

View version information

Knowing the exact version of your management server service or recording server service is an advantage if you need to contact product support.

1. In Management Client's menu bar select *Help* menu, click *About...*.

2. A small dialog opens. Depending on server type, the dialog shows the exact version of your Management Server service or Recording Server service.

Recording server settings

When you configure Recording server settings, specify the following:

Name	Description
Address	IP address (example: 123.123.123.123) or host name (example: ourserver) of the management server to which the recording server should be connected. This information is necessary in order for the recording server to be able to communicate with the management server.
Port	Port number to be used when communicating with the management server. Default is port 9993. You can change this if you need to.
Web server port	Port number to be used for handling web server requests, for example for handling PTZ camera control commands and for browse and live requests from Ocularis Client. Default is port 7563. You can this if you need to.
Alert server port	Port number to be used when the recording server listens for TCP information (some devices use TCP for sending event messages). Default is port 5432. You can change this if you need to.
SMTP server port	Port number to be used when the recording server listens for Simple Mail Transfer Protocol (SMTP) information. Also, some devices use SMTP (e-mail) for sending event messages and/or for sending images to the surveillance system server via e-mail. SMTP is a standard for sending e-mail messages between servers. Default is port 25. You can change this if you need to.
FTP server port	Port number to be used when the recording server listens for FTP information (some devices use FTP for sending event messages. Default is port 21. You can change this if you need to.

Read server service icons - management, recording and failover

The following notification area icons represent the possible states of the Management Server service, Recording Server and Failover Recording Server services. They are all visible on the machines where the service is installed, not in the Management Client (see "Management Client Overview" on page 30):

Manage-ment Server service icon	Recording Server service icon	Failover Recording Server service icon	Description
			**Running.** Reg. failover recording server, it is enabled and started and able to take over from standard recording servers.
			**Stopped.** Reg. failover recording server, it is stopped and no longer taking over from standard recording servers.
			**Starting.** Appears when a server service is in the process of starting. Under normal circumstances, the icon will after a short while change to **Running**.

Manage-ment Server service icon	Recording Server service icon	Failover Recording Server service icon	Description
		**Management and Recording Server service only**	**Stopping.** Appears when a server service is in the process of stopping. Under normal circumstances, the icon will after a short while change to **Stopped**.
Recording Server service only		**Recording Server service only**	**In indeterminate state.** Appears when the **Recording Server service** is initially loaded and until the first information is received, upon which the icon will, under normal circumstances, change to **Starting**, and subsequently to the **Running**.
			**Running offline.** Typically appears when the Recording Server or Failover recording service is running but the Management Server service is not. Reg. **failover recording server**, it typically appears if: • the failover recording server is not enabled (see "About failover recording servers—regular and hot standby" on page 234) through the Management Client. • the failover recording server's information about the management server address is incorrect (see "Change the management server address" on page 241). • the user account under which the Failover Recording Server service runs has no access to your system. To fix this, make sure that the user account specified during installation of the failover recording server, under which the Failover Server service runs, has access to your system with administrator rights. To verify this, do the following: 1. In the Management Client's Site Navigation pane (see "Panes Overview" on page 33), expand ***Security*** and select ***Roles***. 2. In the Overview pane (see "Panes Overview" on page 33)'s roles list, select the *Administrators* role. 3. In the Properties pane's role settings list, check that the required user is listed. 4. If no, add the required user to the *Administrators* role by clicking ***Add...***. Also refer to Work with users, groups and roles (on page 186).

Manage-ment Server service icon	Recording Server service icon	Failover Recording Server service icon	Description
		**Recording Server service only**	**Must be authorized by administrator.** Appears when the Recording Server service is loaded for the first time. Administrators authorize the recording server through the Management Client: In the Management Client's Site Navigation pane, expand the **Servers** list, select the **Recording Server** node then in the Overview pane right-click the required recording server and select **Authorize Recording Server**.

Virus scanning

Virus scanning information

In some cases, OnSSI recommends that you avoid virus scanning, if this is allowed in your organization.

If you use virus scanning software on:

- recording data in databases on recording servers
- data being archived in archiving (see "About storage and archiving" on page 61) locations

It most uses a considerable amount of system resources on scanning.

This may affect system performance negatively, notably scanning of data in databases containing recordings. Some virus scanning software may also temporarily lock each file it scans, which may further impact system performance negatively. Virus scanning may even corrupt recording databases, and render your surveillance system recordings useless.

Therefore:

- Do not use virus scanning on recording server directories containing recording databases (by default *C:\MediaDatabase*\ and all folders under that location, but note that your organization may have specified different recording paths).
- Do not use virus scanning on archiving locations.
- Do not use virus scanning on files with the following file extensions (which are all surveillance system-related):
    - .blk
    - .idx
    - .pic
    - .pqz
    - .sts
    - .ts
- Do not use virus scanning on the management server.

Your organization may have strict guidelines regarding virus scanning, but it is important that the mentioned locations and files are exempt from virus scanning. If allowed, you should disable any virus scanning of recording servers' databases, of any archiving locations as well as on the management server. Consult your organization’s IT system administrator if in doubt.

Tray icon

The following issues are relevant to tasks accessible from the tray icon:

Change Software License Code (on page 49)

Restore system configuration (from manual back up) (on page 229)

Select shared backup folder (on page 229)

Update SQL server address (on page 243)

SNMP

About SNMP support

Depending on the recording component, functionality described here may be limited or unavailable.

Your system supports Simple Network Management Protocol (SNMP), a standard protocol for monitoring and controlling network devices, for managing their configuration, or collecting statistics, etc.

The system will act as an SNMP agent, which can generate an SNMP trap as a result of a triggered rule. A third party SNMP management console can then receive information about the rule-triggering event, and operators of the SNMP management console can configure their system for further action as required.

The implementation uses Microsoft® Windows® SNMP Service for triggering SNMP traps. The SNMP Service must therefore be installed on recording servers. This will—when the SNMP Service has been configured through its own user interface—enable recording servers to send .mib (Management Information Base) files to the SNMP management console.

Install SNMP service

1. On the required recording servers, open Windows' *Add or Remove Programs* dialog (*Start* > *Control Panel* > *Add or Remove Programs*).
2. In the left side of the *Add or Remove Programs* dialog click *Add/Remove Windows Components*. This opens the *Windows Components* wizard.
3. In the wizard, select the check box next to *Management and Monitoring Tools*, then click *Details...* to open the *Management and Monitoring Tools* dialog.
4. In the *Management and Monitoring Tools* dialog, select the check box next to *Simple Network Management Protocol*, then click *OK*.
5. Back in the *Windows Components* wizard, click *Next* and follow the wizard's further steps.

Configure SNMP service

1. On the required recording servers, select *Start* > *Control Panel* > *Administrative Tools* > *Services*.
2. Double-click the SNMP Service.
3. Select the *Traps* tab.
4. Specify a community name, and click *Add to list*.
5. Select the *Destinations* tab.
6. Click *Add*, and specify the IP address or host name of the server running your third party SNMP management station software.
7. Click *OK*.

Daylight saving time

Daylight saving time

Daylight saving time (DST) is the practice of advancing clocks in order for evenings to have more daylight and mornings to have less. Typically, clocks are adjusted forward one hour sometime during the spring season and adjusted backward sometime during the fall season, therefore the saying *spring forward, fall back*. Note that use of DST varies between countries/regions.

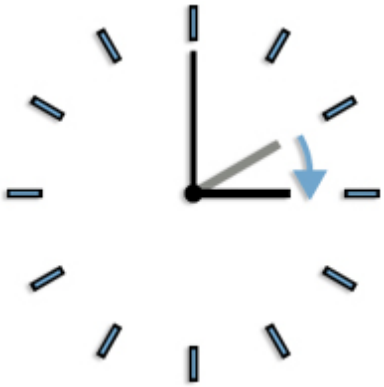

Clocks are adjusted forward when DST starts

When you work with a surveillance system, which is inherently time-sensitive, it is important that you know how the system handles DST.

Spring: Switch from standard time to DST

The change from standard time to DST is not much of an issue since you jump one hour forward. Typically, the clock jumps forward from 02:00 standard time to 03:00 DST, and so the day has 23 hours. In that case, there is no data between 02:00 and 03:00 in the morning since that hour, for that day, did not exist.

Fall: Switch from DST to standard time

When you switch from DST to standard time in the fall, you jump one hour back. Typically, the clock jumps backward from 02:00 DST to 01:00 standard time, repeating that hour, and so the day has 25 hours.

Server-side handling

Your system uses Coordinated Universal Time (UTC), which is the official world reference for time. UTC is not adjusted to reflect switches either to or from DST. Since the system uses UTC, no recordings are ever stored with the same timestamp twice, not even during the DST change hour.

Client-side handling

The client application used for viewing recordings from the system—the Ocularis Client—also uses UTC when displaying recordings. The client takes local time settings (time zone and any DST) from the computer on which the client is used, and converts those time settings to UTC. This means that there is a very simple solution for viewing recordings from the DST change hour.

Viewing DST change hour recordings in clients

When you want to view recordings from the last (most recent) hour of the DST change hour, go ahead and view them.

When you want to view recordings from the first hour of the DST change hour, do the following:

1. On the computer on which the client is used, go to Windows' *Start* menu, and select *Control Panel*.
2. In the Control Panel, double-click *Date and Time*.
3. In the *Date and Time Properties* window, select the *Time Zone* tab.
4. Make sure the *Automatically adjust clock for daylight saving changes* check box is cleared, then click *OK*.

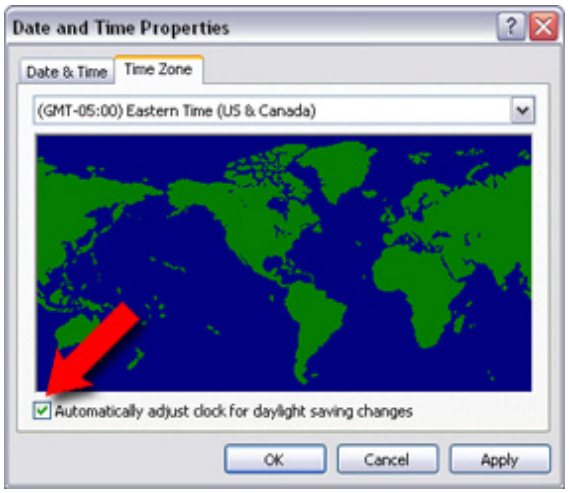

When the Automatically adjust clock for daylight saving changes check box is cleared, recordings from the entire DST period will be Standard Time (or one hour off compared to DST). This means that recordings from the first hour of the DST change hour can now be viewed.

IMPORTANT: When you are done viewing recordings from the first hour of the DST change hour, select the Automatically adjust clock for daylight saving changes check box again to avoid confusion. We recommend not to clear the Automatically adjust clock for daylight saving changes check box unless you specifically need to view recordings from the first hour of the DST change hour.

Multi-domain with one-way trust

Setup with one-way trust

If you run your system in a multi-domain environment, you can configure this setup with one-way trust.

The system is installed on the **trusting** domain and users log in from **trusting** and **trusted** domains.

1. Create a service account in the **trusted** domain. You can name it whatever you want, for example, *svcOnSSI*.
2. Add *svcOnSSI* (example name only) to the following local Windows user groups on the server running the system, in the **trusting** domain:
   - Administrators
   - IIS_IUSRS (Windows Server 2008, necessary for Internet Information Services (IIS) Application Pools)
   - IIS_WPG (Windows Server 2003, necessary for IIS Application Pools).
3. Ensure that the *svcOnSSI* (example name only) account has system administrator rights on your SQL Database or SQL Server Express, either directly or through the BUILTIN\Administrators group.
4. Set the identity of the *ManagementServerAppPool* Application Pool in the IIS to the *svcOnSSI* (example name only) account.
5. Reboot the server to ensure all group membership and permission changes take effect.

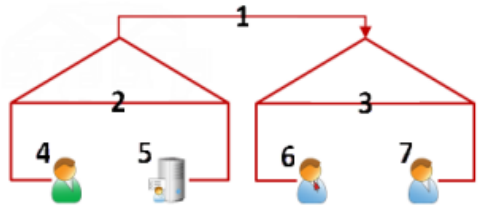

Example illustration of multi-domain environments with one-way trust.

Legend:

1. 1. One-way outgoing domain trust
2. 2. MyDomain.local
3. 3. OtherDomain.edu
4. 4. Trusting domain user
5. 5. Management server
6. 6. OnSSI service account
7. 7. Trusted domain user

IMPORTANT: To add **trusted** domain users to new or existing system roles, log in to Windows as a **trusted** domain user. Next, launch the Management Client and log in as user of either the **trusting** domain or the **trusted** domain. If you log in to Windows as a **trusting** domain user, you are asked for credentials for the **trusted** domain in order to browse for users.

Appendix

Ports used by the system

If nothing else is stated, ports are both inbound and outbound.

- **Port 20 and 21:** Used by **recording servers** to listen for File Transfer Protocol (FTP) information; some devices use FTP for sending event messages. FTP is a standard for exchanging files across networks.
- **Port 25:** Used by **recording servers** to listen for Simple Mail Transfer Protocol (SMTP) information. Also, some devices use SMTP (e-mail) for sending event messages and /or for sending images to the surveillance system server via e-mail. SMTP is a standard for sending e-mail messages between servers.
- **Port 80:** While not directly used by the system, but by **management servers**, port 80 is typically used by the Internet Information Services (IIS) Default Web Site for running the Management Server service.
- **Port 443:** Used by the basic user authentication process where the **management server** must keep this port open at all times.
- **Port 554:** Used by **recording servers** for RTSP traffic which is used for controlling streaming from cameras.
- **Port 1024 and above** (outbound only (except ports listed in the following)): Used by **recording servers** for HTTP traffic between cameras and servers.
- **Port 5210:** Used for communication between **recording servers** and **failover recording servers** when databases are merged after a failover recording server has been running.
- **Port 5432:** Used by **recording servers** to listen for Transmission Control Protocol (TCP) information; some devices use TCP for sending event messages.
- **Port 7563:** Used by **recording servers** and **Ocularis Clients.** The main entry to the recording server where the Image Server interface is implemented. Also used for handling PTZ camera control commands and for retrieving image stream from clients etc.
- **Port 7609:** Used by the **Data Collector Server service** and must always be keep open on the machine running the **Data Collector**.
- **Port 8080**: Used for communication between internal processes on the **management server** only.
- **Port 8844:** Used for User Datagram Protocol (UDP) communication between **failover recording servers**.
- **Port 9993:** Used for communication between **recording servers** and **management servers**.
- **Port 11000:** Used by **failover recording servers** for polling (i.e. regularly checking) the state of **recording servers**.
- **Port 12345:** Used by **management servers** and **Ocularis Client** for communication.
- **Port 65101:** Used between processes on the same machine only – i.e. Inter Process Communication (IPC) on a single machine only.

Index

S